

Provably Robust Score-Based Diffusion Posterior Sampling for Plug-and-Play Image Reconstruction

Xingyu Xu*
Carnegie Mellon University

Yuejie Chi*
Carnegie Mellon University

March 2024; Revised May 2024

Abstract

In a great number of tasks in science and engineering, the goal is to infer an unknown image from a small number of noisy measurements collected from a known forward model describing certain sensing or imaging modality. Due to resource constraints, this image reconstruction task is often extremely ill-posed, which necessitates the adoption of expressive prior information to regularize the solution space. Score-based diffusion models, thanks to its impressive empirical success, have emerged as an appealing candidate of an expressive prior in image reconstruction. In order to accommodate diverse tasks at once, it is of great interest to develop efficient, consistent and robust algorithms that incorporate *unconditional* score functions of an image prior distribution in conjunction with flexible choices of forward models.

This work develops an algorithmic framework for employing score-based diffusion models as an expressive data prior in nonlinear inverse problems with general forward models. Motivated by the plug-and-play framework in the imaging community, we introduce a diffusion plug-and-play method (DPnP) that alternatively calls two samplers, a proximal consistency sampler based solely on the likelihood function of the forward model, and a denoising diffusion sampler based solely on the score functions of the image prior. The key insight is that denoising under white Gaussian noise can be solved *rigorously* via both stochastic (i.e., DDPM-type) and deterministic (i.e., DDIM-type) samplers using the same set of score functions trained for generation. We establish both asymptotic and non-asymptotic performance guarantees of DPnP, and provide numerical experiments to illustrate its promise in solving both linear and nonlinear image reconstruction tasks. To the best of our knowledge, DPnP is the first provably-robust posterior sampling method for nonlinear inverse problems using unconditional diffusion priors. Code is available at <https://github.com/xlxu/diffusion-plug-and-play>.

Keywords: score-based generative models, nonlinear inverse problems, posterior sampling, plug-and-play

Contents

1	Introduction	2
1.1	Our contribution	3
1.2	Related works	4
1.3	Paper organization and notation	5
2	Score-based generative models	5
2.1	The forward process and score functions	6
2.2	The reverse process and sampling	7
3	Posterior sampling via diffusion plug-and-play	8
3.1	Posterior sampling for image reconstruction	8
3.2	Key ingredient: score-based denoising posterior sampling	9
3.3	Our algorithm: diffusion plug-and-play	11

*Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213; Emails: {xingyuxu,yuejiechi}@andrew.cmu.edu.

4	Theoretical analysis	15
4.1	Asymptotic consistency	15
4.2	Non-asymptotic error analysis	16
5	Numerical experiments	17
5.1	Inverse problems	17
5.2	Experimental setups	17
5.3	Results	19
6	Discussion	19
A	Score functions of diffusion denoising samplers	24
A.1	Proof of Lemma 1	24
A.2	Proof of Lemma 2	25
B	Discretization with the exponential integrator	26
B.1	General form of the exponential integrator	26
B.2	Discretization of DDS-DDPM	27
B.3	Discretization of DDS-DDIM	28
B.4	Discretization of PCS	29
C	Proof of main theorems	29
C.1	Proof of Theorem 1	29
C.2	Proof of Lemma 4	30
C.3	Proof of Lemma 5	31
C.4	Proof of Lemma 6	31
C.5	Proof of Theorem 2	32
C.6	Proof of Lemma 9	33
D	Additional experiment details	35
D.1	Hyperparameters of competing methods	35
D.2	Additional performance metrics	35

1 Introduction

In a great number of sensing and imaging applications, the paramount goal is to infer an unknown image $x^* \in \mathbb{R}^d$ from a collection of measurements $y \in \mathbb{R}^m$ that are possibly noisy, incomplete, and even nonlinear, given as

$$y = \mathcal{A}(x^*) + \xi,$$

where $\mathcal{A} : \mathbb{R}^d \rightarrow \mathbb{R}^m$ is the measurement operator encapsulating the *forward model* of the sensing modality, and $\xi \in \mathbb{R}^m$ denotes the measurement noise. Examples include restoration tasks such as inpainting, super-resolution, denoising, as well as imaging tasks such as magnetic resonance imaging (Lustig et al., 2007), optical imaging (Shechtman et al., 2015), microscopy imaging (Huang et al., 2017), radar and sonar imaging (Potter et al., 2010), and many more.

Due to sensing and resource constraints, the problem of image reconstruction is often ill-posed, where the desired resolution of the unknown image overwhelms the set of available observations. Consequently, this necessitates the need of incorporating prior information regarding the unknown image to assist the reconstruction process. Over the years, numerous types of prior information have been considered and adopted, from hand-crafted priors such as subspace or sparsity constraints (Donoho, 2006; Candes and Recht, 2012), to data-driven ones prescribed in the form of neural networks (Ulyanov et al., 2018; Bora et al., 2017). These priors can be regarded as some sort of generative models for the unknown image, which postulate the high-dimensional image admits certain parsimonious representation in a low-dimensional data manifold. It is desirable that the generative models are sufficiently expressive to capture the diversity and

structure of the image class of interest, yet nonetheless, still lead to image reconstruction problems that are computationally tractable.

Score-based diffusion models as an image prior. Recent years have seen tremendous progress on generative artificial intelligence (AI), where it is possible to generate new data samples — such as images, audio, text — at unprecedented resolution and scale from a target distribution given training data. Diffusion models, originally proposed by Sohl-Dickstein et al. (2015), are among one of the most successful frameworks, underneath popular content generators such as DALL·E (Ramesh et al., 2022), Stable Diffusion (Rombach et al., 2022), Imagen (Saharia et al., 2022), and many others. Roughly speaking, score-based diffusion models convert noise into samples that resemble those from a target data distribution, by forming the reverse Markov diffusion process only using the score functions of the data contaminated at various noise levels (Song and Ermon, 2019; Ho et al., 2020; Song et al., 2021b, 2020). In particular, Song et al. (2020) developed a unified framework to interpret score-based diffusion models as reversing certain Stochastic Differential Equations (SDE) using either SDE or probability flow Ordinary Differential Equations (ODE), leading to stochastic (i.e., DDPM-type) and deterministic (i.e., DDIM-type) samplers, respectively. While the DDIM-type sampler is more amenable to acceleration, the DDPM-type sampler tends to generate images of higher quality and diversity when running for a large number of steps (Song et al., 2020).

Thanks to the expressive power of score-based diffusion models in generating complex and fine-grained images, they have emerged as a plausible candidate of an expressive prior in image reconstruction (Song et al., 2021a; Chung et al., 2023a; Feng et al., 2023) via the lens of *Bayesian posterior sampling*. To accommodate diverse applications with various image characteristics and imaging modalities, it is desirable to develop *plug-and-play* methods that do not require training from scratch or end-to-end training for every new imaging task. Nonetheless, despite a flurry of recent efforts, existing algorithms either are computationally expensive (Wu et al., 2023; Cardoso et al., 2023), inconsistent (Chung et al., 2023a; Kavar et al., 2022), or confined to linear inverse problems (Cardoso et al., 2023; Dou and Song, 2024). Therefore, a natural question arises:

Can we develop a practical, consistent and robust algorithm that incorporates score-based diffusion models as an image prior with general (possibly nonlinear) forward models?

1.1 Our contribution

This paper provides an affirmative answer to this question, by developing an algorithmic framework to sample from the posterior distribution of images, where score-based diffusion models are employed as an expressive image prior in nonlinear inverse problems with general forward models. Specifically, our contributions are as follows.

- *Diffusion plug-and-play for posterior sampling.* Motivated by the plug-and-play (Venkatakrishnan et al., 2013) framework in the imaging community, we introduce a diffusion plug-and-play method (DPnP) that alternatively calls two samplers, a *proximal consistency sampler* that aims to generate samples that are more consistent with the measurements, and a *denoising diffusion sampler* that focuses on sampling from the posterior distribution of an easier problem — image denoising under white Gaussian noise — to enforce the prior constraint. Our method is *modular*, in the sense that the *proximal consistency sampler* is solely based on the likelihood function of the forward model, and the denoising diffusion sampler is based solely on the score functions of the image prior.
- *Posterior sampling for image denoising.* While the proximal consistency sampler can be borrowed somewhat straightforwardly from existing literature such as the Metropolis-adjusted Langevin algorithm (Roberts and Rosenthal, 1998), the denoising diffusion sampler, on the other hand, has not been addressed in the literature to the best of our knowledge. Our key insight is that this can be solved via both stochastic (i.e., DDPM-type) or deterministic (i.e., DDIM-type) samplers by carefully choosing the forward SDEs and discretizing the resulting reversal SDE or ODE using the exponential integrator (Zhang and Chen, 2022). Importantly, the denoising diffusion samplers use the same set of unconditional score functions for generation, making it readily implementable without additional training.

- *Theoretical guarantees.* We establish both asymptotic and non-asymptotic performance guarantees of the proposed DPnP method. Asymptotically, we verify the correctness of our method by proving that DPnP converges to the conditional distribution of x^* given measurements y , under a suitable choice of annealing schedule and assuming exact unconditional score functions of the image prior. We next establish a non-asymptotic convergence theory of DPnP, where its performance degenerates gracefully with respect to the errors of the samplers, due to, e.g., score estimation errors and limited sampling steps. To the best of our knowledge, this provides the *first provably-robust* method for *nonlinear* inverse problems using unconditional score-based diffusion priors.

We further provide numerical experiments to illustrate its promise in solving both linear and nonlinear image reconstruction tasks, such as image super-resolution and phase retrieval. Due to its plug-and-play nature, we expect it to be of broad interest to a wide variety of inverse problems.

1.2 Related works

Given its interdisciplinary nature, our work sits at the intersection of generative modeling, computational imaging, optimization and sampling. Here, we discuss some works that are most related to ours.

Algorithmic unrolling and plug-and-play image reconstruction. Composite optimization algorithms, which aim to minimize the sum of a measurement fidelity term and a regularization term promoting desirable solution structures, have been the backbone of inverse problem solvers. To unleash the power of deep learning, [Gregor and LeCun \(2010\)](#) advocates the perspective of algorithmic unrolling, which turns an iterative algorithm into concatenations of linear and nonlinear layers like in a neural network. [Venkatakrishnan et al. \(2013\)](#) recognized that the proximal mapping step in many composite optimization algorithms can be regarded as a denoiser or denoising operator with respect to the given prior, and proposed to “plug in” alternative denoisers, in particular state-of-the-art deep learning denoisers, leading to a class of popular algorithms known as plug-and-play methods ([Buzzard et al., 2018](#)); see [Monga et al. \(2021\)](#) for a review.

Regularization by denoising and score matching. [Vincent \(2011\)](#) pointed out a connection between score matching and image denoising, which is a consequence of the Tweedie’s formula ([Efron, 2011](#)). The regularization by denoising (RED) framework ([Romano et al., 2017](#)) follows the plug-and-play framework to minimize a regularized objective function, where the regularizer is defined based on the plug-in image denoiser; [Reehorst and Schniter \(2018\)](#) later clarified that the RED framework can be interpreted as score matching by denoising using the Tweedie’s formula. [Kawar et al. \(2021\)](#) developed a stochastic image denoiser for posterior sampling of image denoising using annealed Langevin dynamics. [Fang et al. \(2024\)](#) provided a framework to learn exact proximal operators for inverse problems.

Plug-and-play posterior sampling. Motivated by the need to characterize the uncertainty, tackling image reconstruction as posterior sampling from a Bayesian perspective is another important approach. Our method is inspired by the plug-and-play framework but takes on a sampling perspective, exploiting the connection between optimization and sampling ([Wibisono, 2018](#)). Along similar lines, [Laumont et al. \(2022\)](#); [Bouman and Buzzard \(2023\)](#) proposed Bayesian counterparts of plug-and-play for posterior sampling, where they leveraged the connection to score matching for sampling from the image prior, but did not consider score-based diffusion models for the image prior, which is a key aspect of ours; see also [Sun et al. \(2023\)](#). [Coeurdoux et al. \(2023\)](#) extended the split Gibbs sampler ([Vono et al., 2019](#)) in the plug-and-play framework, and advocated the use of score-based diffusion models such as DDPM ([Ho et al., 2020](#)) for image denoising based on heuristic observations. In contrast, we rigorously derive the denoising diffusion samplers from first principles, unraveling critical gaps from naïve applications of the generative samplers to denoising, and offer theoretical guarantees on the correctness of our approach.

Score-based diffusion models as image priors. Several representative methods for solving inverse problems using score-based diffusion priors alternates between taking steps along the diffusion process and projecting onto the measurement constraint, e.g., [Chung et al. \(2023a\)](#); [Kawar et al. \(2022\)](#); [Song et al. \(2023a\)](#); [Chung et al. \(2023b\)](#); [Graikos et al. \(2022\)](#); [Song et al. \(2022\)](#). However, these approaches do not

possess asymptotic consistency guarantees. Song et al. (2023b) proposed to use multiple Monte Carlo samples to reduce bias. On the other hand, Cardoso et al. (2023) developed Monte Carlo guided diffusion methods for Bayesian linear inverse problems which tend to be computationally expensive, and Dou and Song (2024) recently introduced a filtering perspective and applied particle filtering. Although asymptotically consistent, these approaches are limited to linear inverse problems. Trippe et al. (2022); Wu et al. (2023) introduced sequential Monte Carlo (SMC) algorithms for conditional sampling using unconditional diffusion models that are asymptotically exact. Mardani et al. (2023) developed a variational perspective that connects to the regularization by denoising framework. Gupta et al. (2024) showed that the worst-case complexity of diffusion posterior sampling can take super-polynomial time regardless of the algorithm in use.

Theory of diffusion models and score matching. A number of recent papers have studied the non-asymptotic convergence rates of popular diffusion samplers, including but not limited to stochastic DDPM-type samplers (Chen et al., 2022, 2023a; Li et al., 2023; Benton et al., 2024; Tang and Zhao, 2024a), deterministic DDIM-type samplers (Li et al., 2023; Chen et al., 2023b,c), and accelerated samplers (Li et al., 2024a,b). In addition, the statistical efficiency of score matching has also been investigated (Koehler et al., 2023; Pabbaraju et al., 2024).

1.3 Paper organization and notation

The rest of the paper is organized as follows. Section 2 provides preliminaries on score-based generative models. Section 3 provides the problem formulation as well as the developed algorithm, followed by its performance analysis in Section 4. Numerical experiments are provided in Section 5 to elaborate the performance of the proposed algorithm. Last but not least, we conclude in Section 6.

Notation. Let p_x denote the probability distribution of x , and $p_x(\cdot|y)$ denotes the conditional distribution of x given y . We use $X \stackrel{(d)}{=} Y$ to denote random variables X and Y are equivalent in distribution. The matrix I_d denotes an identity matrix of dimension d . For two probability distributions with density $p(x)$ and $q(x)$, the total variation distance between them is

$$\text{TV}(p, q) := \int |p(x) - q(x)| dx.$$

The χ^2 -divergence of p to q is

$$\chi^2(p \| q) := \int \frac{(p(x) - q(x))^2}{q(x)} dx.$$

2 Score-based generative models

In this section, we set up the preliminary on diffusion-based generative models, which we will be relying upon to develop our algorithm. We shall define two stochastic processes in \mathbb{R}^d :

- 1) a forward process

$$x_0 \rightarrow x_1 \rightarrow \cdots \rightarrow x_T$$

that starts with samples from the target image distribution and diffuses into a noise distribution (e.g., standard Gaussians) by gradually injecting noise into the samples;

- 2) a reverse process

$$x_T^{\text{rev}} \rightarrow x_{T-1}^{\text{rev}} \rightarrow \cdots \rightarrow x_0^{\text{rev}}$$

that starts from pure noise (e.g., standard Gaussians) and converts it into samples whose distribution is close to the target image distribution.

2.1 The forward process and score functions

Consider the forward Markov process in \mathbb{R}^d that starts with a sample from the data distribution p_X , and adds noise over the trajectory according to

$$x_0 \sim p^*, \quad (1a)$$

$$x_t = \sqrt{1 - \beta_t} x_{t-1} + \sqrt{\beta_t} w_t, \quad 1 \leq t \leq T, \quad (1b)$$

where $\{w_t\}_{1 \leq t \leq T}$'s are independent standard Gaussian vectors, i.e., $w_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d)$, and $\{\beta_t \in (0, 1)\}$ describes the noise-injection rates used in each step. Therefore, we can write x_t equivalently as

$$x_t := \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, I_d), \quad t = 0, 1, \dots, T. \quad (2)$$

Here, $(\bar{\alpha}_t)_{t=0,1,\dots,T}$ is the *schedule* of diffusion given by

$$\alpha_t := 1 - \beta_t, \quad \bar{\alpha}_t := \prod_{k=1}^t \alpha_k, \quad 1 \leq t \leq T. \quad (3)$$

Clearly, it verifies that $1 \geq \bar{\alpha}_0 > \bar{\alpha}_1 > \dots > \bar{\alpha}_T > 0$. As long as $\bar{\alpha}_T$ is vanishing, it is easy to observe that the distribution of x_T approaches $\mathcal{N}(0, I_d)$.

Score functions. As will be seen, in order to sample from p^* , it turns out to be sufficient to learn the score functions of p_{x_t} at each step of the forward process, defined as

$$s_t^*(x) = \nabla \log p_{x_t}(x), \quad t = 0, 1, \dots, T. \quad (4)$$

An enlightening property (Vincent, 2011) of the score function is that it can be interpreted as the minimum mean-squared error (MMSE) estimate of ε_t given $x_t = x$, fueled by Tweedie's formula:

$$s_t^*(x) = -\frac{1}{\sqrt{1 - \bar{\alpha}_t}} \underbrace{\mathbb{E}_{x_0 \sim p^*, \varepsilon_t \sim \mathcal{N}(0, I_d)}(\varepsilon_t \mid \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \varepsilon_t = x)}_{=: \varepsilon_t^*(x)}. \quad (5)$$

Consequently, this makes it possible to estimate the score functions via learning to denoise (Hyvärinen, 2005), by estimating the denoising function $\varepsilon_t^*(\cdot)$, as typically done in practice (Ho et al., 2020).

Continuous-time perspective. To facilitate understanding, it will be convenient to formulate the diffusion process in continuous time. To distinguish from the discrete-time setting, we use capitalized letters like X to denote the continuous-time diffusion process, and τ to denote continuous time parameter. The continuous-time forward diffusion follows the Ornstein-Uhlenbeck (OU) process¹, defined by the Stochastic Differential Equation (SDE) (Song et al., 2021b):

$$dX_\tau = -X_\tau d\tau + \sqrt{2} dB_\tau, \quad \tau \geq 0, \quad X_0 \sim p^*, \quad (6)$$

where $(B_\tau)_{\tau \geq 0}$ is the standard d -dimensional Brownian motion. It can be shown that (Doob, 1942; Evans, 2012) the marginal distribution of X_τ for $\tau \geq 0$ is

$$X_\tau \stackrel{(d)}{=} e^{-\tau} X_0 + \sqrt{1 - e^{-2\tau}} \varepsilon, \quad X_0 \sim p^*, \quad \varepsilon \sim \mathcal{N}(0, I_d). \quad (7)$$

Comparing (2) and (7), it can be checked that the discrete-time diffusion process can be embedded into the continuous-time one via the time change $t \mapsto \frac{1}{2} \log \frac{1}{\bar{\alpha}_t}$, in the sense that $x_t^* \stackrel{(d)}{=} X_{\frac{1}{2} \log \frac{1}{\bar{\alpha}_t}}$.

Similar to the discrete-time case, the continuous-time score function is defined by

$$s^{\text{cont}}(\tau, x) = \nabla \log p_{X_\tau}(x) \quad (8)$$

¹In the literature, other processes such as Variance-Exploding SDE (VE-SDE) are also used. Our theory also applies to these processes with straightforward modifications.

$$= -\frac{1}{\sqrt{1-e^{-2\tau}}} \underbrace{\mathbb{E}_{X_0 \sim p^*, \varepsilon \sim \mathcal{N}(0, I_d)}(\varepsilon | e^{-\tau} X_0 + \sqrt{1-e^{-2\tau}} \varepsilon = x)}_{=: \varepsilon^{\text{cont}}(\tau, x)},$$

where the second line follows again from the Tweedie's formula, where $\varepsilon^{\text{cont}}(\tau, x)$ is the continuous-time counterpart of $\varepsilon_t^*(\cdot)$. By the change-of-time argument mentioned above, we have the following correspondence between the discrete-time and the continuous-time score functions:

$$s^{\text{cont}}\left(\frac{1}{2} \log \frac{1}{\hat{\alpha}_t}, x\right) = s_t^*(x). \quad (9)$$

2.2 The reverse process and sampling

To enable sampling, one needs to “reverse” the forward diffusion process. Fortunately, it is possible to leverage classical theory (Anderson, 1982; Ambrosio et al., 2005) to reverse the SDE, and apply discretization to the time-reversal processes to collect samples. We shall describe two popular approaches below, corresponding to stochastic (i.e., DDPM-type) and deterministic (i.e., DDIM-type) samplers respectively following primarily the framework set forth in Song et al. (2021b).

Time-reversed SDEs and probability flow ODEs. Let us begin with the more general theory of *reversing* SDEs, which will be useful in future sections. Consider a SDE given by

$$dM_\tau = f(M_\tau)d\tau + \sqrt{\beta}dB_\tau, \quad \tau \geq 0, \quad M_0 \sim p_{M_0}, \quad (10)$$

where $f: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a deterministic function and $\beta > 0$ is a constant. For any positive time $\tau_\infty > 0$, define the reversed time parameter

$$\tau^{\text{rev}} := \tau^{\text{rev}}(\tau) = \tau_\infty - \tau. \quad (11)$$

We are now ready to describe the time-reversed processes.

- 1) The *time-reversed SDE* of (10) on the time interval $[0, \tau_\infty]$ is defined as

$$dM_{\tau^{\text{rev}}}^{\text{rev}} = (-f(M_{\tau^{\text{rev}}}^{\text{rev}}) + \beta \nabla \log p_{M_{\tau^{\text{rev}}}^{\text{rev}}}(M_{\tau^{\text{rev}}}^{\text{rev}})) d\tau + \sqrt{\beta} d\tilde{B}_\tau, \quad \tau \in [0, \tau_\infty], \quad M_{\tau_\infty}^{\text{rev}} \sim p_{M_{\tau_\infty}}, \quad (12)$$

where \tilde{B} is an independent copy of B , i.e., another Brownian motion. It is a classical result in stochastic analysis (Anderson, 1982) that the reversed process M^{rev} shares the same path distribution as G , i.e.,

$$(M_\tau^{\text{rev}})_{\tau \in [0, \tau_\infty]} \stackrel{(d)}{=} (M_\tau)_{\tau \in [0, \tau_\infty]}.$$

In other words, the joint distribution of $(M_{\tau_1}^{\text{rev}}, M_{\tau_2}^{\text{rev}}, \dots, M_{\tau_k}^{\text{rev}})$ for any $0 \leq \tau_1 \leq \tau_2 \leq \dots \leq \tau_k \leq \tau_\infty$, for any integer $k \geq 1$, coincides with that of $(M_{\tau_1}, M_{\tau_2}, \dots, M_{\tau_k})$.

- 2) In place of the reversed SDE in (12), it is possible to consider the following probability flow ODE (Ambrosio et al., 2005; Song et al., 2021b):

$$dM_{\tau^{\text{rev}}}^{\text{rev}} = \left(-f(M_{\tau^{\text{rev}}}^{\text{rev}}) + \frac{\beta}{2} \nabla \log p_{M_{\tau^{\text{rev}}}^{\text{rev}}}(\tau^{\text{rev}}, M_{\tau^{\text{rev}}}^{\text{rev}}) \right) d\tau, \quad \tau \in [0, \tau_\infty], \quad M_{\tau_\infty}^{\text{rev}} \sim p_{M_{\tau_\infty}}. \quad (13)$$

The reversed ODE satisfies a slightly weaker guarantee than that of the reversed SDE, which nevertheless suffices for most practical purposes (Song et al., 2021b):

$$M_\tau^{\text{rev}} \stackrel{(d)}{=} M_\tau, \quad \tau \in [0, \tau_\infty].$$

Note that the reversed ODE only guarantees identical marginal distribution for each M_τ^{rev} , whereas the reversed SDE guarantees identical joint distribution.

Specializing the above to the OU process (6) with proper discretization then leads to popular samplers used for generation, as follows.

DDPM-type stochastic samplers. Specializing the time-reversed SDE (12) to the OU process gives

$$dX_{\tau^{\text{rev}}}^{\text{rev}} = (X_{\tau^{\text{rev}}}^{\text{rev}} + 2s^{\text{cont}}(\tau^{\text{rev}}, X_{\tau^{\text{rev}}}^{\text{rev}}))d\tau + \sqrt{2}d\tilde{B}_\tau, \quad \tau \in [0, \tau_\infty], \quad X_{\tau_\infty}^{\text{rev}} \sim p_{X_{\tau_\infty}}.$$

As $\tau_\infty \rightarrow \infty$, it can be seen from (7) that $p_{X_{\tau_\infty}}$ converges in distribution to $\varepsilon \sim \mathcal{N}(0, I_d)$. Thus the solution of the above SDE can be approximated by initializing $X_{\tau_\infty}^{\text{rev}} \sim \mathcal{N}(0, I_d)$ instead. The DDPM sampler (Ho et al., 2020) can be viewed as a discretization of this SDE (Song et al., 2021b).

DDIM-type deterministic samplers. On the other hand, the probability flow ODE (13) for the OU process reads as

$$dX_{\tau^{\text{rev}}}^{\text{rev}} = (X_{\tau^{\text{rev}}}^{\text{rev}} + s^{\text{cont}}(\tau^{\text{rev}}, X_{\tau^{\text{rev}}}^{\text{rev}}))d\tau, \quad \tau \in [0, \tau_\infty], \quad X_{\tau_\infty}^{\text{rev}} \sim p_{X_{\tau_\infty}}. \quad (14)$$

Again, as $\tau_\infty \rightarrow \infty$, one may approximate the initialization with $X_{\tau_\infty}^{\text{rev}} \sim \mathcal{N}(0, I_d)$. It is known that the popular DDIM sampler (Song et al., 2021b, 2020) is a discretization of this ODE (Zhang and Chen, 2022). The ODE-based deterministic samplers allow more aggressive choice of discretization schedules, as well as fast ODE solvers (Lu et al., 2022), enabling significantly accelerated sampling process compared to the SDE-based stochastic samplers.

Remark 1. In correspondence with the discrete-time schedule, one might take $\tau_\infty = -\frac{1}{2} \log \bar{\alpha}_T$. We refer interested readers to Song et al. (2021b); Tang and Zhao (2024b) for more details.

3 Posterior sampling via diffusion plug-and-play

3.1 Posterior sampling for image reconstruction

We are interested in solving (possibly nonlinear) inverse problems, where the aim is to infer an unknown image $x^* \in \mathbb{R}^d$ from its measurements $y \in \mathbb{R}^m$,² given by

$$y = \mathcal{A}(x^*) + \xi,$$

where $\mathcal{A} : \mathbb{R}^d \rightarrow \mathbb{R}^m$ is the measurement operator underneath the forward model, and ξ denotes measurement noise. It has been well-understood that *prior information* of x^* plays an important role in solving inverse problems that are otherwise ill-posed, enabling successful reconstruction with much less measurements and higher accuracy. At the same time, it is desirable to understand and quantify the uncertainty in image reconstruction, especially when the available measurements are rather limited.

Posterior sampling. In this work, we focus on the Bayesian setting where the prior information of x^* is provided in the form of some prior distribution $p^*(\cdot)$, i.e.,

$$x^* \sim p^*(x), \quad (15)$$

The *posterior distribution* given measurements y is defined as

$$p^*(x|y) \propto p^*(x) p(y|x^* = x) = p^*(x) e^{\mathcal{L}(x;y)}. \quad (16)$$

Here, $\mathcal{L}(\cdot; y)$ is the log-likelihood function of the measurements. For example, when the noise $\xi \sim \mathcal{N}(0, \sigma^2 I_m)$ is standard Gaussian, it follows that

$$\mathcal{L}(x; y) = -\frac{1}{2\sigma^2} \|y - \mathcal{A}(x)\|^2 - \frac{m}{2} \log(2\pi\sigma^2).$$

Notwithstanding, our framework allows flexible choices of the forward model and the noise distributions. In addition, while this formulation is derived from probabilistic interpretations, it also subsumes the “reward-guided” or “loss-guided” setting (Song et al., 2023b), where \mathcal{L} can be viewed as a reward function or a negative loss function, both of which characterize preference over structural properties of x^* . In all these settings, it will be useful to bear in mind the intuition that the higher value of \mathcal{L} corresponds to better consistency with the measurements, higher rewards, etc.

²For simplicity, we limit our presentation to the real-valued case; our framework generalizes to the complex-valued case in a straightforward manner.

Assumption on the forward model. Throughout the paper, for simplicity, we make the following mild assumption on \mathcal{L} , which is applicable to many applications of interest.

Assumption 1. We assume $\mathcal{L}(\cdot; y)$ is differentiable almost everywhere, and $\sup_{x \in \mathbb{R}^d} \mathcal{L}(x; y) < \infty$.

Goal. Our goal is to sample \hat{x} from the posterior distribution

$$\hat{x} \sim p^*(\cdot | y)$$

given estimates $\hat{s}_t(x)$ (resp. $\hat{\varepsilon}_t(x)$) of the *unconditional* score functions $s_t^*(x)$ (resp. the noise function $\varepsilon_t^*(x)$) in (4), assuming knowledge of the likelihood function $\mathcal{L}(\cdot; y)$.

3.2 Key ingredient: score-based denoising posterior sampling

We begin with an inspection on one of the most fundamental inverse problems: denoising under white Gaussian noise. We shall demonstrate how to solve this problem via stochastic (i.e., DDPM-type) and deterministic (i.e., DDIM-type) denoising diffusion samplers using the same set of unconditional score functions trained for generation. As shall be elucidated shortly, the denoising diffusion samplers turn out to be an important building block in our algorithm for general inverse problems.

Image denoising under white Gaussian noise. Suppose that we have access to a noisy version of $x^* \sim p^*$ contaminated by white Gaussian noise, given by

$$x_{\text{noisy}} = x^* + \xi, \quad \xi \sim \mathcal{N}(0, \eta^2 I_d), \quad (17)$$

where $\eta > 0$ is the noise intensity *assumed to be known*. Our goal is to sample from $p^*(\cdot | x_{\text{noisy}})$ given the score estimates $\hat{s}_t(x)$ (resp. the noise estimates $\hat{\varepsilon}_t(x)$). We will develop our score-based denoising posterior sampler, termed DDS, with two variants, DDS-DDPM and DDS-DDIM, which can be viewed as analogues of the well-known DDPM and DDIM samplers in unconditional score-based sampling respectively. Before proceeding, it is worth highlighting that the two variants will be derived from different forward diffusion processes, since we observe the resulting variants empirically lead to more competitive performance.

A stochastic DDPM-type sampler via heat flow. We begin with a stochastic DDPM-type sampler for denoising, termed DDS-DDPM. We divide our development into the following steps.

- 1) *Step 1: introducing the heat flow.* Let us introduce a *heat flow* with initial distribution p^* , defined by the following SDE:

$$dY_\tau = dB_\tau, \quad \tau \geq 0, \quad Y_0 \sim p^*, \quad (18)$$

where $(B_\tau)_{\tau \geq 0}$ is the standard d -dimensional Brownian motion. The solution of (18) is simply

$$Y_\tau = Y_0 + B_\tau, \quad \tau \geq 0. \quad (19)$$

Since $B_\tau \sim \mathcal{N}(0, \tau I_d)$, it readily follows that $B_{\eta^2} \stackrel{(d)}{=} \xi$, which together with $Y_0 \sim p^*$ yield the important observation that $x_{\text{noisy}} = x^* + \xi$ can be viewed as an endpoint of the heat flow, in the sense that

$$x_{\text{noisy}} = x^* + \xi \stackrel{(d)}{=} Y_{\eta^2}.$$

- 2) *Step 2: reversing the heat flow.* Following similar reasonings in Section 2, the next step boils down to reverse the heat flow (18). The time-reversal of the heat flow SDE (18) is (cf. (12)) given by

$$dY_{\eta^2-\tau}^{\text{rev}} = \nabla \log p_{Y_{\eta^2-\tau}}(Y_{\eta^2-\tau}^{\text{rev}}) d\tau + d\tilde{B}_\tau, \quad \tau \in [0, \eta^2], \quad Y_{\eta^2}^{\text{rev}} \sim p_{Y_{\eta^2}}, \quad (20)$$

where $(\tilde{B}_\tau)_{\tau \geq 0}$ is an independent copy of $(B_\tau)_{\tau \geq 0}$. As introduced earlier, the virtue of the time-reversed SDE (20) is that it produces a process Y_τ^{rev} with the same *path* distribution as Y_τ , i.e.,

$$(Y_\tau^{\text{rev}})_{\tau \in [0, \eta^2]} \stackrel{(d)}{=} (Y_\tau)_{\tau \in [0, \eta^2]}.$$

In particular, the joint distribution of $(Y_0^{\text{rev}}, Y_{\eta^2}^{\text{rev}})$ is the same as that of $(Y_0, Y_{\eta^2}) \stackrel{(d)}{=} (x^*, x_{\text{noisy}})$. This implies that the conditional distribution $p^*(\cdot | x_{\text{noisy}})$ is the same as $p_{Y_0^{\text{rev}}}(\cdot | Y_{\eta^2}^{\text{rev}} = x_{\text{noisy}})$. Surprisingly, the latter admits a simple interpretation: $p_{Y_0^{\text{rev}}}(\cdot | Y_{\eta^2}^{\text{rev}} = x_{\text{noisy}})$ is the distribution of Y_0^{rev} when we initialize (20) with $Y_{\eta^2}^{\text{rev}} = x_{\text{noisy}}$! Therefore, sampling the posterior $p^*(\cdot | x_{\text{noisy}})$ amounts to solving the following simple SDE:

$$dY_{\eta^2-\tau}^{\text{rev}} = \nabla \log p_{Y_{\eta^2-\tau}}(Y_{\eta^2-\tau}^{\text{rev}})d\tau + d\tilde{B}_\tau, \quad \tau \in [0, \eta^2], \quad Y_{\eta^2}^{\text{rev}} = x_{\text{noisy}}. \quad (21)$$

3) *Step 3: connecting the score functions.* It is now immediate to arrive at our proposed stochastic sampler DDS-DDPM by discretization of this SDE (21), which requires knowledge of the score functions $\nabla \log p_{Y_\tau}(\cdot)$. A key observation is that they can in fact be computed from the score function $s^{\text{cont}}(\tau, x)$ (cf. (8)), due to the following lemma, whose proof is provided in Appendix A.1.

Lemma 1 (Score function of Y_τ). *For $\tau \geq 0$, we have*

$$\nabla \log p_{Y_\tau}(x) = \frac{1}{\sqrt{1+\tau}} s^{\text{cont}}\left(\frac{1}{2} \log(1+\tau), \frac{x}{\sqrt{1+\tau}}\right).$$

We leave the detailed discretization procedure with an exponential integrator (Zhang and Chen, 2022) to Appendix B.2. The resulting sampler, DDS-DDPM, is summarized in Algorithm 1.

Algorithm 1 Denoising Diffusion Sampler (stochastic) DDS-DDPM($x_{\text{noisy}}, \hat{s}, \eta$)

Input: noisy data $x_{\text{noisy}} \in \mathbb{R}^d$, score estimates $\hat{s} := \{\hat{s}_t(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d, t = 1, \dots, T\}$ or noise estimates $\hat{\varepsilon} = \{\hat{\varepsilon}_t(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d, t = 1, \dots, T\}$, and noise level $\eta > 0$.

Scheduling: Compute the diffusion schedule $(\tau_t)_{0 \leq t \leq T'}$ by

$$\tau_t = \bar{\alpha}_t^{-1} - 1, \quad 0 \leq t \leq T',$$

where

$$T' := \max \left\{ t : 0 \leq t \leq T, \bar{\alpha}_t > \frac{1}{\eta^2 + 1} \right\}.$$

Initialization: Set $\hat{x}_{T'} = x_{\text{noisy}}$.

Diffusion: for $t = T', T' - 1, \dots, 1$ do

$$\hat{x}_{t-1} = \hat{x}_t - 2(\sqrt{\tau_t} - \sqrt{\tau_{t-1}}) \hat{\varepsilon}_t + \sqrt{\tau_t - \tau_{t-1}} w_t, \quad w_t \sim \mathcal{N}(0, I_d).$$

where

$$\hat{\varepsilon}_t := \hat{\varepsilon}_t(\sqrt{\bar{\alpha}_t} \hat{x}_t) = -\frac{1}{\sqrt{1 - \bar{\alpha}_t}} \hat{s}_t(\sqrt{\bar{\alpha}_t} \hat{x}_t).$$

Output: \hat{x}_0 .

A deterministic DDIM-type sampler via OU process. We next develop a deterministic DDIM-type sampler for denoising, termed DDS-DDIM, presented in Algorithm 2.

1) *Step 1: introducing a posterior-initialized OU process.* To sample from the posterior distribution $p^*(\cdot | x_{\text{noisy}})$, we first introduce a random variable w which has (unconditional) distribution

$$p_w(x) := p^*(x^* = x | x^* + \xi = x_{\text{noisy}}), \quad (22)$$

in the same form of the desired posterior distribution $p^*(\cdot|x_{\text{noisy}})$. Here, since the noisy observation x_{noisy} is given, we regard it as fixed.³ We then further introduce $z = w - x_{\text{noisy}}$, which is a “centered” version of w , whose distribution is

$$p_z(x) := p_w(x + x_{\text{noisy}}) = p^*(x^* = x + x_{\text{noisy}} | x^* + \xi = x_{\text{noisy}}).$$

The OU process with initial distribution p_z is defined by the SDE:

$$dZ_\tau = -Z_\tau d\tau + dB_\tau, \quad \tau \geq 0, \quad Z_0 \sim p_z, \quad (23)$$

where B_τ is the standard d -dimensional Brownian motion. As in (7), the marginal distribution of Z_τ is given by

$$Z_\tau \stackrel{(d)}{=} e^{-\tau} Z_0 + \sqrt{1 - e^{-2\tau}} \varepsilon, \quad Z_0 \sim p_z, \quad \varepsilon \sim \mathcal{N}(0, I_d), \quad \tau \geq 0. \quad (24)$$

- 2) *Step 2: reversing the OU process.* Following similar reasonings in Section 2, reversing the OU process (23) will enable us to generate samples $z \sim p_z$. Then we can set $w = z + x_{\text{noisy}}$, which, by definition, has distribution p_w defined in (22), and is a sample from the desired posterior distribution $p^*(\cdot|x_{\text{noisy}})$. We are thus led to solve the time-reversed probability flow ODE (cf. (13)) of (23), given by

$$dZ_{\tau^{\text{rev}}} = (Z_{\tau^{\text{rev}}}^{\text{rev}} + \nabla \log p_{Z_{\tau^{\text{rev}}}}(\tau^{\text{rev}}, Z_{\tau^{\text{rev}}}^{\text{rev}})) d\tau, \quad \tau \in [0, \tau_\infty], \quad Z_{\tau_\infty}^{\text{rev}} \sim \mathcal{N}(0, I_d), \quad \tau^{\text{rev}} = \tau_\infty - \tau. \quad (25)$$

- 3) *Step 3: connecting the score functions.* We are now one step away from our proposed deterministic sampler DDS-DDIM, which is derived by discretization of the ODE (25). We need to know the score functions $\nabla \log p_{Z_\tau}(\cdot)$, which again can be computed from the score function $s^{\text{cont}}(\tau, x)$ (cf. (8)), as documented by the following lemma, whose proof is provided in Appendix A.2.

Lemma 2 (Score function of Z_τ). *For $\tau \geq 0$, we have*

$$\nabla \log p_{Z_\tau}(x) = -\frac{e^{2\tau} x}{\eta^2 + e^{2\tau} - 1} + \frac{e^{\tau - \tilde{\tau}} \eta^2}{\eta^2 + e^{2\tau} - 1} s^{\text{cont}}\left(\tilde{\tau}, e^{-\tilde{\tau}} x_{\text{noisy}} + \frac{e^{\tau - \tilde{\tau}} \eta^2 x}{\eta^2 + e^{2\tau} - 1}\right), \quad (26)$$

where

$$\tilde{\tau} := \tilde{\tau}(\tau) = \frac{1}{2} \log \left(\frac{\eta^2 (e^{2\tau} - 1)}{\eta^2 + e^{2\tau} - 1} + 1 \right). \quad (27)$$

After plugging this into (25) and solving the ODE for Z_τ^{rev} , we see that $Z_0^{\text{rev}} + x_{\text{noisy}}$ is the desired sample from the posterior distribution $p^*(\cdot|x_{\text{noisy}})$, as argued before. Numerically, the ODE (25) is solved by discretization with an exponential integrator (Zhang and Chen, 2022), resulting in the sampler DDS-DDIM as summarized in Algorithm 2. Details of the discretization, which are somewhat nuanced, are left to Appendix B.3.

3.3 Our algorithm: diffusion plug-and-play

Now we turn to the general setting where the measurement operator \mathcal{A} is arbitrary. From the factorization of posterior distribution in (16), one intuitively understands that a posterior sampler must obey two constraints simultaneously: (i) the *data prior constraint*, corresponding to the first factor $p^*(x)$, which imposes that the posterior sampler should be less likely to sample at those points where $p^*(x)$ is small; (ii) the *measurement consistency constraint*, corresponding to the second factor $e^{\mathcal{L}(x;y)}$, which imposes that $\mathcal{A}(x) \approx y$.

³Technically, this can be done by conditioning on x_{noisy} throughout our discussion of DDS-DDIM.

Algorithm 2 Denoising Diffusion Sampler (deterministic) DDS-DDIM($x_{\text{noisy}}, \hat{s}, \eta$)

Input: noisy data $x_{\text{noisy}} \in \mathbb{R}^d$, score estimates $\hat{s} := \{\hat{s}_t(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d, t = 1, \dots, T\}$ or noise estimates $\hat{\varepsilon} = \{\hat{\varepsilon}_t(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^d, t = 1, \dots, T\}$, and noise level $\eta > 0$.

Scheduling: Compute the diffusion schedule $(\bar{u}_t)_{0 \leq t \leq T'}$ by

$$\bar{u}_t = \frac{(\eta^2 + 1)\bar{\alpha}_t - 1}{\eta^2 + \bar{\alpha}_t - 1}, \quad 0 \leq t \leq T',$$

where

$$T' := \max \left\{ t : 0 \leq t \leq T, \bar{\alpha}_t > \frac{1}{\eta^2 + 1} \right\}.$$

Initialization: Draw $z_{T'} \sim \mathcal{N}(0, I_d)$.

Diffusion: for $t = T', T' - 1, \dots, 1$ do

$$z_{t-1} = \frac{\sqrt{(\eta^2 - 1)\bar{u}_{t-1} + 1}}{\sqrt{(\eta^2 - 1)\bar{u}_t + 1}} z_t + \sqrt{(\eta^2 - 1)\bar{u}_{t-1} + 1} \cdot (h(\eta, \bar{u}_{t-1}) - h(\eta, \bar{u}_t)) \hat{\varepsilon}_t,$$

where

$$h(\eta, u) := -\arctan \frac{\eta}{\sqrt{u^{-1} - 1}},$$

$$\hat{\varepsilon}_t := \hat{\varepsilon}_t \left(\sqrt{\bar{\alpha}_t} x_{\text{noisy}} + \frac{\eta^2 \sqrt{\bar{u}_t \bar{\alpha}_t} z_t}{(\eta^2 - 1)\bar{u}_t + 1} \right) = -\frac{1}{\sqrt{1 - \bar{\alpha}_t}} \hat{s}_t \left(\sqrt{\bar{\alpha}_t} x_{\text{noisy}} + \frac{\eta^2 \sqrt{\bar{u}_t \bar{\alpha}_t} z_t}{(\eta^2 - 1)\bar{u}_t + 1} \right).$$

Output: $x_{\text{noisy}} + z_0$.

A prelude: proximal gradient method. Informally speaking, bridging the perspective of optimization and sampling (Wibisono, 2018), posterior sampling can be viewed as a “soft” solution to the following optimization problem:

$$\max_{x \in \mathbb{R}^d} \mathcal{L}(x; y) + \log p^*(x). \quad (28)$$

Instead of producing the point estimate that maximizes $\mathcal{L}(x; y) + \log p^*(x)$, posterior sampling produces samples from the posterior distribution $p^*(x) e^{\mathcal{L}(x; y)} = e^{\mathcal{L}(x; y) + \log p^*(x)}$ instead, which allows characterizing the underlying uncertainty. For better understanding, it is useful to bear in mind the special case when the image prior p^* is supported on some low-dimensional manifold⁴ \mathcal{M} . In this setting, we notice that $\log p^*(x) = -\infty$ for $x \notin \mathcal{M}$, hence the optimization problem (28) is implicitly constrained in $x \in \mathcal{M}$.

Recall the well-known proximal gradient method (Parikh and Boyd, 2014) for solving (28), where one initializes a random $\hat{x}_0 \in \mathbb{R}^d$ and uses the following update rule

$$\hat{x}_{k+1} = \text{Prox}_{\mathcal{M}, \eta_k} (\hat{x}_k + \eta_k \nabla_{\hat{x}_k} \mathcal{L}(\hat{x}_k; y)), \quad k = 0, 1, \dots$$

Here, $\eta_k > 0$ is the stepsize at the k -th iteration, and $\text{Prox}_{\mathcal{M}, \eta} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the proximal operator defined by

$$\text{Prox}_{\mathcal{M}, \eta}(x) := \underset{x' \in \mathbb{R}^d}{\text{argmin}} -\log p^*(x') + \frac{1}{2\eta^2} \|x' - x\|^2. \quad (29)$$

Intuitively, one may view $\text{Prox}_{\mathcal{M}, \eta}(x)$ as some kind of denoising to make x more consistent with its structural property (Venkatakrisnan et al., 2013). The proximal gradient method alternatively applies two operations:

- (i) *Gradient step to enforce the measurement consistency constraint.* The gradient step tries to boost consistency $y \approx \mathcal{A}(x)$ via moving along the direction to increase $\mathcal{L}(\cdot; y)$.

⁴This assumption, known as the manifold hypothesis, is commonly adopted as a flexible structural characterization of high-dimensional data. We mention it here to facilitate the understanding of our design, which will not be imposed in our algorithm or analysis.

Algorithm 3 Diffusion Plug-and-Play (DPnP)

Input: Measurements $y \in \mathbb{R}^m$, log-likelihood function $\mathcal{L}(\cdot; y)$ of the forward model, score estimates $\hat{s}(\cdot)$, annealing schedule $(\eta_k)_{0 \leq k \leq K}$.

Initialization: Sample $\hat{x}_0 \sim \mathcal{N}(0, \frac{\eta_0}{4} I_d)$

Alternating sampling: for $k = 0, 1, 2, \dots, K - 1$ do

(1) *Proximal consistency sampler:* Sample $\hat{x}_{k+\frac{1}{2}} \propto \exp\left(\mathcal{L}(\cdot; y) - \frac{1}{2\eta_k^2} \|\cdot - \hat{x}_k\|^2\right)$ using subroutine PCS($\hat{x}_k, y, \mathcal{L}, \eta_k$).

(2) *Denoising diffusion sampler:* Sample $\hat{x}_{k+1} \sim \exp\left(\log p^*(x) - \frac{1}{2\eta_k^2} \|x - \hat{x}_{k+\frac{1}{2}}\|^2\right)$ using subroutine DDS-DDPM($\hat{x}_{k+\frac{1}{2}}, \hat{s}, \eta_k$) or DDS-DDIM($\hat{x}_{k+\frac{1}{2}}, \hat{s}, \eta_k$).

Output: \hat{x}_K .

- (ii) *Proximal mapping to enforce the data prior constraint.* The proximal step moves the iterate towards those points that increase $\log p^*(x)$. In particular, when p^* is supported on a low-dimensional manifold \mathcal{M} , the proximal map forces x to reside in \mathcal{M} .

Diffusion plug-and-play (DPnP). Although the proximal gradient method does not apply to the posterior sampling problem directly, we will apply the idea of alternatively enforcing these two constraints from a sampling perspective in the same spirit of Vono et al. (2019); Lee et al. (2021); Bouman and Buzzard (2023). Our algorithm, dubbed diffusion plug-and-play (DPnP), alternates between two samplers, the denoising diffusion sampler (DDS) and the proximal consistency sampler (PCS), which can be viewed as the substitutes for the proximal operator and the gradient step respectively. Given the iterate \hat{x}_k and annealing parameter η_k at the k -th iteration, DPnP proceeds with the following two steps:

- (i) *Proximal consistency sampler to enforce the measurement consistency constraint.* DPnP draws a sample $\hat{x}_{k+\frac{1}{2}}$ from the distribution proportional to

$$\exp\left(\mathcal{L}(x; y) - \frac{1}{2\eta_k^2} \|x - \hat{x}_k\|^2\right)$$

to promote the image to be consistent with the measurements. This step, which we denote as the *proximal consistency sampler*, can be achieved by small modifications of standard algorithms such as Metropolis-Adjusted Langevin Algorithm (MALA) (Roberts and Rosenthal, 1998) given in Algorithm 4.⁵

- (ii) *Denoising diffusion sampler to enforce the data prior constraint.* DPnP next draws a sample \hat{x}_{k+1} from the distribution proportional to

$$\exp\left(-\left(-\log p^*(x) + \frac{1}{2\eta_k^2} \|x - \hat{x}_{k+\frac{1}{2}}\|^2\right)\right) \propto p^*(x) e^{-\frac{1}{2\eta_k^2} \|x - \hat{x}_{k+\frac{1}{2}}\|^2} \propto p^*(x^* = x | x^* + \eta_k w = \hat{x}_{k+\frac{1}{2}}) \quad (30)$$

to promote the image to be consistent with the prior, where $w \sim \mathcal{N}(0, I_d)$. The last step, which follows from the Bayes' rule, makes it clear that this step can be precisely achieved by the denoising diffusion sampler (developed in Section 3.2) using solely the unconditional score function, with two options given in Algorithm 1 and Algorithm 2.

Combining both steps lead to the proposed DPnP method described in Algorithm 3.

Discussions of DPnP. Some comments about the proposed DPnP method are in order.

⁵For practical reasons, we apply the exponential integrator to tame the discretization error of the drift term. For a derivation of this algorithm as discretization of the continuous-time Langevin dynamics, please refer to Appendix B.4.

Algorithm 4 Proximal Consistency Sampler $\text{PCS}(x, y, \mathcal{L}, \eta)$ (adapted from Metropolis-Adjusted Langevin Algorithm (Roberts and Rosenthal, 1998))

Input: starting point $x \in \mathbb{R}^d$, measurements $y \in \mathbb{R}^m$, log-likelihood function of the forward model $\mathcal{L}(\cdot; y)$, proximal parameter $\eta > 0$.

Hyperparameter: Langevin stepsize γ , and the number of iterations N .

Initialization: $z_0 = x$.

Update: for $n = 0, 1, \dots, N - 1$ do

(1) **One step of discretized Langevin:** Set $r = e^{-\gamma/\eta^2}$, and

$$z_{n+\frac{1}{2}} = rz_n + (1-r)x + \eta^2(1-r)\nabla_{z_n}\mathcal{L}(z_n; y) + \eta\sqrt{1-r^2}w_n, \quad w_n \sim \mathcal{N}(0, I_d).$$

This is equivalent to drawing $z_{n+\frac{1}{2}}$ from a distribution with density $Q(\cdot; z_n)$, where

$$Q(z'; z) = \frac{1}{(2\pi(1-r^2))^{d/2}} \exp\left(-\frac{\|z' - (rz + (1-r)x + \eta^2(1-r)\nabla_z\mathcal{L}(z; y))\|^2}{2(1-r^2)}\right).$$

(2) **Metropolis adjustment:** Compute

$$q = \frac{\exp\left(\mathcal{L}(z_{n+\frac{1}{2}}; y) - \frac{1}{2\eta^2}\|z_{n+\frac{1}{2}} - x\|^2\right)}{\exp\left(\mathcal{L}(z_n; y) - \frac{1}{2\eta^2}\|z_n - x\|^2\right)} \cdot \frac{Q(z_n; z_{n+\frac{1}{2}})}{Q(z_{n+\frac{1}{2}}; z_n)},$$

and set

$$z_{n+1} = \begin{cases} z_{n+\frac{1}{2}}, & \text{with probability } \min(1, q), \\ z_n & \text{with probability } 1 - \min(1, q). \end{cases}$$

Output: z_N .

- The proximal consistency sampler PCS can be viewed as a “soft” version of the proximal point method (Drusvyatskiy, 2017). This can be seen from a first-order approximation: the maximum likelihood of the distribution $\exp\left(\mathcal{L}(\cdot; y) - \frac{1}{2\eta_k^2}\|\cdot - \hat{x}_k\|^2\right)$ is attained at the point $x' \in \mathbb{R}^d$ satisfying

$$\nabla_{x'}\mathcal{L}(x'; y) - \frac{1}{\eta_k^2}(x' - \hat{x}_k) = 0, \quad \implies \quad x' = \hat{x}_k + \eta_k^2\nabla_{x'}\mathcal{L}(x'; y) \approx \hat{x}_k + \eta_k^2\nabla_{\hat{x}_k}\mathcal{L}(\hat{x}_k; y),$$

Therefore, the proximal consistency sampler draws random samples “concentrated” around x' , which approximates the implicit proximal point update, akin to a gradient step at \hat{x}_k .

- On the other end, the denoising posterior sampler DDS can be regarded as a “soft” version of the proximal operator by comparing (30) with (29). In particular, when p^* is supported on a low-dimensional manifold \mathcal{M} , it forces x to reside in \mathcal{M} , like the proximal map. To see this, note that denoising posterior distribution vanishes outside \mathcal{M} by (30). This is in contrast with existing algorithms (Romano et al., 2017; Chung et al., 2023a; Laumont et al., 2022; Bouman and Buzzard, 2023; Sun et al., 2023), which use the conditional expectation map $x_{\text{noisy}} \mapsto \mathbb{E}_{x^* \sim p^*, \varepsilon \sim \mathcal{N}(0, I_d)}(x^*|x^* + \eta\varepsilon = x_{\text{noisy}})$ (the stepsize $\eta > 0$ may vary with time) (in view of (5)) in lieu of the proximal operator for denoising. Given the presumption that the conditional expectation is a weighted average of points on \mathcal{M} , we see that when the manifold \mathcal{M} is nonlinear, this map may generate points that are far apart from \mathcal{M} , which violates the data prior constraint and hence degrades reconstruction quality.
- We note that the proximal consistency sampler PCS in fact admits a simple form in the special case where the forward model \mathcal{A} is linear, i.e. $\mathcal{A}(x) = Ax$ for some matrix $A \in \mathbb{R}^{m \times d}$ (assume A is of full row rank), and the measurement noise $\xi \sim \mathcal{N}(0, \Sigma)$ is Gaussian. In this situation, the log-likelihood

function becomes

$$\mathcal{L}(x; y) = -\frac{1}{2}(y - Ax)^\top \Sigma^{-1}(y - Ax) + \text{constant}.$$

Then the target distribution of PCS can be rewritten as a Gaussian distribution with proper mean and covariance, namely,

$$\exp\left(\mathcal{L}(\cdot; y) - \frac{1}{2\eta_k^2} \|\cdot - \hat{x}_k\|^2\right) \propto \exp\left(-\frac{1}{2}(\cdot - \tilde{x}_k)^\top \tilde{\Sigma}_k^{-1}(\cdot - \tilde{x}_k)\right) = \mathcal{N}(\tilde{x}_k, \tilde{\Sigma}_k),$$

where

$$\tilde{x}_k = \left(A^\top \Sigma^{-1} A + \frac{1}{\eta_k^2} I_d\right)^{-1} \left(A^\top \Sigma^{-1} y + \frac{1}{\eta_k^2} \hat{x}_k\right), \quad \tilde{\Sigma}_k = \left(A^\top \Sigma^{-1} A + \frac{1}{\eta_k^2} I_d\right)^{-1}.$$

Therefore, the proximal consistency sampler PCS can be implemented directly by

$$\hat{x}_{k+\frac{1}{2}} = \text{PCS}(\hat{x}_k, y, \mathcal{L}, \eta_k) = \tilde{x}_k + \tilde{\Sigma}_k^{1/2} w_k, \quad w_k \sim \mathcal{N}(0, I_d).$$

4 Theoretical analysis

In this section, we establish both asymptotic and non-asymptotic performance guarantees of DPnP, whose proofs are postponed to Appendix C.

4.1 Asymptotic consistency

We first collect the asymptotic correctness of our subroutines PCS and DDS in the following two lemmas. The correctness of PCS is actually well-known, see e.g., Tierney (1994, Corollary 2).

Lemma 3 (Correctness of PCS). *Under Assumption 1, with notation in Algorithm 4, in the continuous-time limit:*

$$\gamma \rightarrow 0, \quad N \rightarrow \infty,$$

the algorithm PCS outputs samples with distribution $\propto \exp(\mathcal{L}(\cdot; y) + \frac{1}{2\eta} \|\cdot - x\|^2)$.

The next lemma guarantees the correctness of DDS with exact unconditional score functions.

Lemma 4 (Correctness of DDS). *Assume the score function estimation \hat{s}_t is accurate, i.e. $\hat{s}_t = s_t^*$. In the continuous-time limit:*

$$T \rightarrow \infty, \quad \bar{\alpha}_T \rightarrow 0, \quad \frac{\bar{\alpha}_{t-1}}{\bar{\alpha}_t} \rightarrow 1, \quad \text{uniformly in } t,$$

both DDS-DDIM and DDS-DDPM output samples x obeying the denoising posterior distribution $p^(x^* = x | x^* + \eta\varepsilon = x_{\text{noisy}})$, $\varepsilon \sim \mathcal{N}(0, I_d)$.*

We are now ready to state our main result, which concerns the asymptotic correctness of DPnP.

Theorem 1 (Asymptotic consistency of DPnP). *Under the settings of Lemma 4 and Lemma 3, the following holds. Let $\varepsilon_1 > \varepsilon_2 > \dots$ be a decreasing sequence of positive numbers satisfying $\lim_{l \rightarrow \infty} \varepsilon_l = 0$, and $0 = k_0 < k_1 < k_2 < \dots$ be an increasing sequence of integers. Set the annealing schedule as follows:*

$$\eta_k = \varepsilon_l, \quad \text{for } k_{l-1} \leq k < k_l, \quad l = 1, 2, \dots$$

Let $\min_{l'=1,2,\dots} |k_{l'} - k_{l'-1}| \rightarrow \infty$, the output \hat{x}_{k_l} of DPnP converges in distribution to the posterior distribution $p^(\cdot|y)$ for $l \rightarrow \infty$.*

In words, Theorem 1 establishes the asymptotic consistency of DPnP under fairly mild assumptions on the forward model (cf. Assumption 1): as long as the sampled distributions of DDS and PCS are exact, then running DPnP with a slowly diminishing annealing schedule of $\{\eta_k\}$ will output samples approaching the desired posterior distribution $p^*(\cdot|y)$ when the number of iterations K goes to infinity.

4.2 Non-asymptotic error analysis

We now step away from the idealized setting when the sampled distributions of DDS and PCS are exact. In practice, there are many sources of errors that can influence the sampled distributions of DDS and PCS: non-diminishing γ and finite number of sampling steps N in PCS, and score estimation error and finite number of discretization steps T in DDS. In effect, these non-idealities will make the subroutines PCS and DDS *inexact*. In other words, the distribution they generate will slightly deviate from the distribution they ought to sample from. In this paper, we model such deviations by the *total variation* distance from the distribution generated by PCS (resp. DDS) to the ideal distribution proportional to $\exp(\mathcal{L}(x; y) - \frac{1}{2\eta_k} \|x - \hat{x}_k\|^2)$ (resp. $p^*(x^* = x | x^* + \eta_k \varepsilon = \hat{x}_{k+\frac{1}{2}})$) uniformly over all iterations. Analyzing these total variations errors is out of the scope of this paper, and we point the interested readers to parallel lines of works, e.g., [Li et al. \(2023\)](#); [Mangoubi and Vishnoi \(2019\)](#); [Chewi et al. \(2021\)](#), among many others. In our analysis, we will assume a black-box bound for the total variation errors of PCS and DDS, which can be combined with existing analyses of the respective samplers to bound the iteration complexity of DPnP.

Theorem 2 (Non-asymptotic robustness of DPnP). *With the notation in DPnP (Algorithm 3), set $\eta_k \equiv \eta > 0$. Under Assumption 1, there exists $\lambda := \lambda(p^*, \mathcal{L}, \eta) \in (0, 1)$, such that the following holds. Define a stationary distribution π_η by*

$$\pi_\eta(x) \propto p^*(x) q_\eta(x),$$

where q_η is defined by

$$q_\eta(x) := e^{\mathcal{L}(\cdot; y)} * p_{\eta\varepsilon}(x) = \frac{1}{(2\pi)^{d/2}\eta^d} \int e^{\mathcal{L}(x'; y) - \frac{1}{2\eta^2} \|x - x'\|^2} dx', \quad \varepsilon \sim \mathcal{N}(0, I_d), \quad (31)$$

where $*$ denotes convolution. If PCS has error at most ε_{PCS} in total variation and DDS has error at most ε_{DDS} in total variation per iteration, then for any accuracy goal $\varepsilon_{\text{acc}} > 0$, with $K \asymp \frac{\log(1/\varepsilon_{\text{acc}})}{1-\lambda}$, we have

$$\text{TV}(p_{\hat{x}_K}, \pi_\eta) \lesssim \varepsilon_{\text{acc}} \sqrt{\chi^2(p_{\hat{x}_1} \| \pi_\eta)} + \frac{1}{1-\lambda} (\varepsilon_{\text{DDS}} + \varepsilon_{\text{PCS}}) \log \left(\frac{1}{\varepsilon_{\text{acc}}} \right). \quad (32)$$

Before interpreting Theorem 2, we observe that $q_0(x) = e^{\mathcal{L}(x; y)}$, thus $\pi_0(x) \propto p^*(x) e^{\mathcal{L}(x; y)}$ coincides with the desired posterior distribution $p^*(\cdot | y)$. Thus Theorem 2 tells us that, assuming a constant annealing schedule $\eta_k = \eta$, the output of DPnP converges in total variation to the distribution π_η , which is a distorted version of the desired posterior distribution up to level η , with sufficiently many iterations.

A few remarks are in order.

Non-diminishing η . It can be seen from Theorem 2 that even with a nonzero η , DPnP already enforces the data prior strictly. On the other hand, the measurement consistency is distorted by an order of η . This is usually tolerable, since the measurements are themselves contaminated by noise, thus when η is smaller than the noise level, the distortion would be tolerable. In practice, it is beneficial to choose an annealing schedule of $\{\eta_k\}$, which will be elaborated in Section 5.

Spectral gap and worst-case convergence rate. The term $1 - \lambda$ is known as the *spectral gap* of the associated Markov chain of DPnP. In many situations, it can be shown that $1 - \lambda \gtrsim \frac{\eta}{\text{poly}(d)}$, see e.g. [Vono et al. \(2022\)](#) for the case of log-concave p^* and negative quadratic \mathcal{L} . In such cases, the factor $\frac{1}{1-\lambda}$ in the right hand side of (32) can also be improved significantly ([Altschuler and Chewi, 2023](#)). However, under the minimal assumption in this paper and without any additional assumption on p^* and \mathcal{L} , $1 - \lambda$ can be exponentially small in the worst case, thus our result does not contradict the worst-case lower bound in [Gupta et al. \(2024\)](#).

Provable robustness. Theorem 2 indicates the performance of DPnP degenerates gracefully in the presence of sampling errors. To the best of our knowledge, this is the first provably consistent and robust posterior sampling method for nonlinear inverse problems using score-based diffusion priors.

5 Numerical experiments

We provide numerical evidence to corroborate the promise of DPnP in solving both linear and nonlinear image reconstruction tasks. We denote DPnP with the subroutines DDS-DDPM and DDS-DDS-DDIM as DPnP-DDPM and DPnP-DDS-DDIM respectively.

5.1 Inverse problems

We consider the following linear and nonlinear inverse problems in our experiments.

Phase retrieval. We consider phase retrieval with a coded mask, which is a classical inverse problem (Candes et al., 2015). For a 256×256 image x (for each color channel) in our experiments, we first generate a random mask $M \in \mathbb{R}^{256 \times 256}$ (which is shared across color channels), then apply Fourier transform \mathcal{F} to $M \odot x$, where \odot denotes the Hadamard (entrywise) product, and finally preserve only the magnitudes of the Fourier transform. Formally, the forward measurement operator is $\mathcal{A}(x) = \text{mag}(\mathcal{F}(M \odot x))$, where $\text{mag}(\cdot)$ computes the entrywise magnitude of a matrix with complex entries. The measurement noise is again set to be white Gaussian, with variance 0.2.

Quantized sensing. Quantized sensing refers to the task of reconstructing an image from its low-bit quantized version. Here, the forward measurement operator is a one-bit per channel, dithered quantization operator. More precisely, it applies entrywise the following stochastic function Q with dithering level $\theta > 0$:

$$Q(\text{pixel}) = \begin{cases} 1, & \text{with probability } \frac{e^{\text{pixel}/\theta}}{1+e^{\text{pixel}/\theta}} \\ -1, & \text{with probability } \frac{1}{1+e^{\text{pixel}/\theta}}, \end{cases}$$

where $\text{pixel} \in [-1, 1]$ is the value of each pixel in each channel. The measurements in quantized sensing are therefore one-bit-per-channel images. The dithering level θ is set to 0.4 in our experiments.

Super resolution. The forward model for super-resolution is the bicubic downsampling operator (Keys, 1981), which is a linear operator (in fact, a block Hankel matrix). We use a downsampling ratio of 4 in all our experiments. The measurement noise is set to be white Gaussian, with variance 0.2.

5.2 Experimental setups

We compare DPnP with the state-of-the-art DPS algorithm (Chung et al., 2023a) and LGD-MC algorithm (Song et al., 2023b) on the FFHQ validation dataset (Karras et al., 2019) and the ImageNet validation dataset (Russakovsky et al., 2015). We use the same pre-trained score functions as in Chung et al. (2023a),⁶ and all images are normalized to fit into the range $[-1, 1]$.

Annealing schedule. For DPnP, we use a heuristic strategy to choose the annealing schedule η_k in DPnP (Algorithm 3). As seen in the theoretical analysis (Theorem 2), if we set all the $\eta_k \equiv \eta$ for some constant $\eta > 0$, then DPnP converges to a distribution π_η , which can be regarded as a version of the posterior distribution $p^*(\cdot|y)$ distorted by an order of $O(\eta)$. The smaller η is, the more accurate the final distribution will be. On the other hand, it was also seen that in many cases, the spectral gap is $\Omega(\eta)$, hence the convergence time is $O(\frac{1}{\eta})$. Therefore, smaller η would make it take longer to converge.⁷


To strike a balance between the accuracy and the convergence rate, we adapt an gradually decreasing schedule for η_k , similar to Bouman and Buzzard (2023). In the first few iterations, we set η_k to be a large constant. After this initial phase, we decrease η_k slowly, eventually to η_N which is chosen to be a small constant. An example of such an annealing schedule is

$$\eta_0 = \eta_1 = \dots = \eta_{K_0}, \quad \eta_0 > 0 \text{ is a large constant,}$$


⁶<https://github.com/DPS2022/diffusion-posterior-sampling>

⁷Strictly speaking, while the number of iterations required to converge increases as η gets smaller, the computational complexity of PCS and DDS per iteration will decrease. However, in experiments, we observed that the latter effect is not strong enough to compensate for the increase in overall complexity caused by the former.


Figure 1: Samples of different algorithms for phase retrieval, quantized sensing, and super resolution, where DPnP generate images of higher quality and recover fine details of the image more faithfully than the state-of-the-art DPS (Chung et al., 2023a) and LGD-MC (Song et al., 2023b) algorithms.



(a) phase retrieval



(b) quantized sensing



(c) super resolution

$$\eta_k = (\eta_K/\eta_0)^{\frac{k-K_0}{K-K_0}} \eta_0, \quad K_0 < k \leq K, \quad \eta_K > 0 \text{ a small constant},$$

where $K_0 < K$ is the length of the initial phase, which can be chosen as, e.g., $K_0 = K/5$. For all the numerical experiments, we set $\eta_0 = 0.4$, $\eta_N = 0.15$, $K_0 = 4$, $K = 20$. The annealing schedule $\{\eta_k\}$ of DPnP is fixed across *all* tasks, while DPS and LGD-MC are fine-tuned with reasonable effort for best performance (see Appendix D.1 for the choice of hyperparameters). All experiments are run on a single Nvidia L40 GPU.

Initialization. In Algorithm 3, the initial guess \hat{x}_0 is set to be a properly scaled Gaussian random vector. Notwithstanding, from Theorem 2 it can be inferred that using a heuristic posterior sampler as the initializer could decrease $\chi^2(p_{\hat{x}_1} \parallel \pi_\eta)$, hence potentially improve the convergence speed of DPnP. By using existing algorithms like DPS or LGD-MC as initializers, DPnP can improve upon the results of existing algorithms towards the correct posterior distribution efficiently and provably. In our experiments, we find it helpful to initialize DPnP with LGD-MC, which accelerates the algorithm significantly.

5.3 Results

Visual results. The samples generated by different algorithms are shown in Figure 1. It can be seen from these results that, DPnP is capable of solving both linear and nonlinear problems, and, in comparison with state-of-the-art algorithms, performs better in recovering fine and crisper details.

Performance metric. We report the performance metric of DPnP in terms of LPIPS and PSNR — which are two of the most relevant metrics for inverse problems — on the FFHQ and ImageNet datasets in Table 1 and Table 2, respectively. Since DPnP-DDIM has similar performance with DPnP-DDPM but admits much faster implementation, only DPnP-DDIM is evaluated. We refer to comparisons to other metrics such as FID and SSIM to Appendix D.2. It can be seen that DPnP-DDIM performs strongly on both datasets, albeit taking about 1.5x more computation time.

Table 1: Evaluation of solving inverse problems on FFHQ 256×256 validation dataset (1k samples).

Algorithm	Super-resolution (4x, linear)		Phase retrieval (nonlinear)		Quantized sensing (nonlinear)		Time per sample
	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	
DPnP-DDIM (ours)	0.301	24.2	0.376	22.4	0.293	24.2	~ 90s
DPS (Chung et al., 2023a)	0.331	23.1	0.490	17.4	0.367	21.7	~ 60s
LGD-MC ($n = 5$) (Song et al., 2023b)	0.318	23.9	0.522	16.4	0.317	23.9	~ 60s

Table 2: Evaluation of solving inverse problems on ImageNet 256×256 validation dataset (1k samples).

Algorithm	Super-resolution (4x, linear)		Phase retrieval (nonlinear)		Quantized sensing (nonlinear)		Time per sample
	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	LPIPS ↓	PSNR ↑	
DPnP-DDIM (ours)	0.416	21.6	0.562	13.4	0.363	23.0	~ 240s
DPS (Chung et al., 2023a)	0.473	20.2	0.677	13.4	0.542	18.7	~ 150s
LGD-MC ($n = 5$) (Song et al., 2023b)	0.416	20.9	0.592	12.8	0.384	22.3	~ 150s

6 Discussion

This paper sets forth a rigorous and versatile algorithmic framework called DPnP for solving nonlinear inverse problems via posterior sampling, using image priors prescribed by score-based diffusion models with general forward models. DPnP alternates between two sampling steps implemented by DDS and PCS, to promote consistency with the data prior constraint and the measurement constraint respectively. We provide both asymptotic and non-asymptotic convergence guarantees, establishing DPnP as the first provably consistent and robust score-based diffusion posterior sampling method for general nonlinear inverse problems. Our work opens up many interesting questions, which we single out a few below.

- *Refined theory.* It is possible to further develop end-to-end finite-time convergence theory for DPnP, and obtain refined theory for structured inverse problems with additional properties of the image prior and the forward model.

- *Accelerated posterior sampling.* Due to the modular design, it is straightforward to incorporate existing accelerated samplers for both DDS (Lu et al., 2022) and PCS (Ma et al., 2021) to speed up the inference, which is of broad practical interest.
- *Non-differentiable forward models.* While we assume the log-likelihood function $\mathcal{L}(\cdot; y)$ to be differentiable to apply MALA for PCS, it is straightforward to adopt other samplers that only assume zero-order oracle access to $\mathcal{L}(\cdot; y)$ for non-differentiable forward models.
- *Guided generation.* While we focus on solving inverse problems, our design might provide some insights into improving the quality of controlled or guided generation (Song et al., 2023b) as well.

We leave these directions to future work.

Acknowledgements

This work is supported in part by Office of Naval Research under N00014-19-1-2404, and by National Science Foundation under DMS-2134080 and ECCS-2126634. X. Xu is also gratefully supported by the Axel Berny Presidential Graduate Fellowship at Carnegie Mellon University.

References

- Altschuler, J. M. and Chewi, S. (2023). Faster high-accuracy log-concave sampling via algorithmic warm starts. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2169–2176. IEEE.
- Ambrosio, L., Gigli, N., and Savaré, G. (2005). *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media.
- Anderson, B. D. (1982). Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326.
- Benton, J., De Bortoli, V., Doucet, A., and Deligiannidis, G. (2024). Nearly d -linear convergence bounds for diffusion models via stochastic localization. In *The Twelfth International Conference on Learning Representations*.
- Bora, A., Jalal, A., Price, E., and Dimakis, A. G. (2017). Compressed sensing using generative models. In *International conference on machine learning*, pages 537–546. PMLR.
- Bouman, C. A. and Buzzard, G. T. (2023). Generative plug and play: Posterior sampling for inverse problems. *arXiv preprint arXiv:2306.07233*.
- Bourbaki, N. (2023a). *Théories spectrales Chapitres 1 et 2*. Springer.
- Bourbaki, N. (2023b). *Théories spectrales: Chapitres 3 à 5*. Springer Nature.
- Buzzard, G. T., Chan, S. H., Sreehari, S., and Bouman, C. A. (2018). Plug-and-play unplugged: Optimization-free reconstruction using consensus equilibrium. *SIAM Journal on Imaging Sciences*, 11(3):2001–2020.
- Candes, E. and Recht, B. (2012). Exact matrix completion via convex optimization. *Communications of the ACM*, 55(6):111–119.
- Candes, E. J., Li, X., and Soltanolkotabi, M. (2015). Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007.
- Cardoso, G., Idrissi, Y. J. E., Corff, S. L., and Moulines, E. (2023). Monte carlo guided diffusion for Bayesian linear inverse problems. *arXiv preprint arXiv:2308.07983*.

- Chen, H., Lee, H., and Lu, J. (2023a). Improved analysis of score-based generative modeling: User-friendly bounds under minimal smoothness assumptions. In *International Conference on Machine Learning*, pages 4735–4763.
- Chen, S., Chewi, S., Lee, H., Li, Y., Lu, J., and Salim, A. (2023b). The probability flow ODE is provably fast. *Neural Information Processing Systems*.
- Chen, S., Chewi, S., Li, J., Li, Y., Salim, A., and Zhang, A. R. (2022). Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. *arXiv preprint arXiv:2209.11215*.
- Chen, S., Daras, G., and Dimakis, A. (2023c). Restoration-degradation beyond linear diffusions: A non-asymptotic analysis for DDIM-type samplers. In *International Conference on Machine Learning*, pages 4462–4484.
- Chewi, S., Lu, C., Ahn, K., Cheng, X., Le Gouic, T., and Rigollet, P. (2021). Optimal dimension dependence of the metropolis-adjusted langevin algorithm. In *Conference on Learning Theory*, pages 1260–1300. PMLR.
- Chung, H., Kim, J., McCann, M. T., Klasky, M. L., and Ye, J. C. (2023a). Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations*.
- Chung, H., Lee, S., and Ye, J. C. (2023b). Fast diffusion sampler for inverse problems by geometric decomposition. *arXiv preprint arXiv:2303.05754*.
- Coeurdoux, F., Dobigeon, N., and Chainais, P. (2023). Plug-and-play split Gibbs sampler: embedding deep generative priors in Bayesian inference. *arXiv preprint arXiv:2304.11134*.
- Donoho, D. L. (2006). Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306.
- Doob, J. L. (1942). The Brownian movement and stochastic equations. *Annals of Mathematics*, 43(2):351–369.
- Dou, Z. and Song, Y. (2024). Diffusion posterior sampling for linear inverse problem solving: A filtering perspective. In *International Conference on Learning Representations*.
- Drusvyatskiy, D. (2017). The proximal point method revisited. *arXiv preprint arXiv:1712.06038*.
- Efron, B. (2011). Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614.
- Evans, L. C. (2012). *An introduction to stochastic differential equations*, volume 82. American Mathematical Soc.
- Fang, Z., Buchanan, S., and Sulam, J. (2024). What’s in a prior? Learned proximal networks for inverse problems. In *The Twelfth International Conference on Learning Representations*.
- Feng, B. T., Smith, J., Rubinstein, M., Chang, H., Bouman, K. L., and Freeman, W. T. (2023). Score-based diffusion models as principled priors for inverse imaging. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10486–10497.
- Graikos, A., Malkin, N., Jovic, N., and Samaras, D. (2022). Diffusion models as plug-and-play priors. *Advances in Neural Information Processing Systems*, 35:14715–14728.
- Gregor, K. and LeCun, Y. (2010). Learning fast approximations of sparse coding. In *Proceedings of the 27th international conference on international conference on machine learning*, pages 399–406.
- Gupta, S., Jalal, A., Parulekar, A., Price, E., and Xun, Z. (2024). Diffusion posterior sampling is computationally intractable. *arXiv preprint arXiv:2402.12727*.
- Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851.

- Huang, J., Sun, M., Ma, J., and Chi, Y. (2017). Super-resolution image reconstruction for high-density three-dimensional single-molecule microscopy. *IEEE Transactions on Computational Imaging*, 3(4):763–773.
- Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4).
- Karras, T., Laine, S., and Aila, T. (2019). A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410.
- Kawar, B., Elad, M., Ermon, S., and Song, J. (2022). Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606.
- Kawar, B., Vaksman, G., and Elad, M. (2021). Stochastic image denoising by sampling from the posterior distribution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1866–1875.
- Keys, R. (1981). Cubic convolution interpolation for digital image processing. *IEEE transactions on acoustics, speech, and signal processing*, 29(6):1153–1160.
- Koehler, F., Heckett, A., and Risteski, A. (2023). Statistical efficiency of score matching: The view from isoperimetry. *International Conference on Learning Representations*.
- Laumont, R., Bortoli, V. D., Almansa, A., Delon, J., Durmus, A., and Pereyra, M. (2022). Bayesian imaging using plug & play priors: when Langevin meets Tweedie. *SIAM Journal on Imaging Sciences*, 15(2):701–737.
- Lee, Y. T., Shen, R., and Tian, K. (2021). Structured logconcave sampling with a restricted Gaussian oracle. In *Conference on Learning Theory*, pages 2993–3050. PMLR.
- Li, G., Huang, Y., Efimov, T., Wei, Y., Chi, Y., and Chen, Y. (2024a). Accelerating convergence of score-based diffusion models, provably. *arXiv preprint arXiv:2403.03852*.
- Li, G., Huang, Z., and Wei, Y. (2024b). Towards a mathematical theory for consistency training in diffusion models. *arXiv preprint arXiv:2402.07802*.
- Li, G., Wei, Y., Chen, Y., and Chi, Y. (2023). Towards faster non-asymptotic convergence for diffusion-based generative models. *arXiv preprint arXiv:2306.09251*.
- Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., and Zhu, J. (2022). DPM-Solver: A fast ODE solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787.
- Lustig, M., Donoho, D., and Pauly, J. M. (2007). Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 58(6):1182–1195.
- Ma, Y.-A., Chatterji, N. S., Cheng, X., Flammarion, N., Bartlett, P. L., and Jordan, M. I. (2021). Is there an analog of Nesterov acceleration for gradient-based MCMC?
- Mangoubi, O. and Vishnoi, N. K. (2019). Nonconvex sampling with the metropolis-adjusted langevin algorithm. In *Conference on learning theory*, pages 2259–2293. PMLR.
- Mardani, M., Song, J., Kautz, J., and Vahdat, A. (2023). A variational perspective on solving inverse problems with diffusion models. *arXiv preprint arXiv:2305.04391*.
- Monga, V., Li, Y., and Eldar, Y. C. (2021). Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44.

- Pabbaraju, C., Rohatgi, D., Sevekari, A. P., Lee, H., Moitra, A., and Risteski, A. (2024). Provable benefits of score matching. *Advances in Neural Information Processing Systems*, 36.
- Parikh, N. and Boyd, S. (2014). Proximal algorithms. *Foundations and trends® in Optimization*, 1(3):127–239.
- Polyanskiy, Y. and Wu, Y. (2024). *Information theory: From coding to learning*. Cambridge university press.
- Potter, L. C., Ertin, E., Parker, J. T., and Cetin, M. (2010). Sparsity and compressed sensing in radar imaging. *Proceedings of the IEEE*, 98(6):1006–1020.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. (2022). Hierarchical text-conditional image generation with CLIP latents. *arXiv preprint arXiv:2204.06125*.
- Reehorst, E. T. and Schniter, P. (2018). Regularization by denoising: Clarifications and new interpretations. *IEEE transactions on computational imaging*, 5(1):52–67.
- Roberts, G. O. and Rosenthal, J. S. (1998). Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268.
- Romano, Y., Elad, M., and Milanfar, P. (2017). The little engine that could: Regularization by denoising (RED). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252.
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E. L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al. (2022). Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494.
- Saloff-Coste, L. (1997). *Lectures on finite Markov chains*, pages 301–413. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Schaefer, H. (2012). *Banach Lattices and Positive Operators*, volume 215. Springer Science & Business Media.
- Shechtman, Y., Eldar, Y. C., Cohen, O., Chapman, H. N., Miao, J., and Segev, M. (2015). Phase retrieval with application to optical imaging: a contemporary overview. *IEEE signal processing magazine*, 32(3):87–109.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. (2015). Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265.
- Song, B., Kwon, S. M., Zhang, Z., Hu, X., Qu, Q., and Shen, L. (2023a). Solving inverse problems with latent diffusion models via hard data consistency. *arXiv preprint arXiv:2307.08123*.
- Song, J., Meng, C., and Ermon, S. (2020). Denoising diffusion implicit models. In *International Conference on Learning Representations*.
- Song, J., Vahdat, A., Mardani, M., and Kautz, J. (2022). Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations*.
- Song, J., Zhang, Q., Yin, H., Mardani, M., Liu, M.-Y., Kautz, J., Chen, Y., and Vahdat, A. (2023b). Loss-guided diffusion models for plug-and-play controllable generation. In *International Conference on Machine Learning*, pages 32483–32498. PMLR.

- Song, Y. and Ermon, S. (2019). Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32.
- Song, Y., Shen, L., Xing, L., and Ermon, S. (2021a). Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations*.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. (2021b). Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations*.
- Sun, Y., Wu, Z., Chen, Y., Feng, B. T., and Bouman, K. L. (2023). Provable probabilistic imaging using score-based generative priors. *arXiv preprint arXiv:2310.10835*.
- Tang, W. and Zhao, H. (2024a). Contractive diffusion probabilistic models. *arXiv preprint arXiv:2401.13115*.
- Tang, W. and Zhao, H. (2024b). Score-based diffusion models via stochastic differential equations—a technical tutorial. *arXiv preprint arXiv:2402.07487*.
- Tierney, L. (1994). Markov Chains for Exploring Posterior Distributions. *The Annals of Statistics*, 22(4):1701–1728.
- Trippe, B. L., Yim, J., Tischer, D., Baker, D., Broderick, T., Barzilay, R., and Jaakkola, T. (2022). Diffusion probabilistic modeling of protein backbones in 3d for the motif-scaffolding problem. *arXiv preprint arXiv:2206.04119*.
- Ulyanov, D., Vedaldi, A., and Lempitsky, V. (2018). Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454.
- Venkatakrishnan, S. V., Bouman, C. A., and Wohlberg, B. (2013). Plug-and-play priors for model based reconstruction. In *IEEE Global Conference on Signal and Information Processing*, pages 945–948. IEEE.
- Vincent, P. (2011). A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674.
- Vono, M., Dobigeon, N., and Chainais, P. (2019). Split-and-augmented Gibbs sampler-application to large-scale inference problems. *IEEE Transactions on Signal Processing*, 67(6):1648–1661.
- Vono, M., Paulin, D., and Doucet, A. (2022). Efficient MCMC sampling with dimension-free convergence rate using ADMM-type splitting. *Journal of Machine Learning Research*, 23(25):1–69.
- Wibisono, A. (2018). Sampling as optimization in the space of measures: The Langevin dynamics as a composite optimization problem. In *Conference on Learning Theory*, pages 2093–3027. PMLR.
- Wu, L., Trippe, B., Naesseth, C., Blei, D., and Cunningham, J. P. (2023). Practical and asymptotically exact conditional sampling in diffusion models. *Advances in Neural Information Processing Systems*, 36.
- Zhang, Q. and Chen, Y. (2022). Fast sampling of diffusion models with exponential integrator. In *The Eleventh International Conference on Learning Representations*.

A Score functions of diffusion denoising samplers

A.1 Proof of Lemma 1

Proof. The marginal distribution (19) of the heat flow can be written as

$$Y_\tau \stackrel{(d)}{=} Y_0 + \sqrt{\tau}\varepsilon, \quad Y_0 \sim p^*, \quad \varepsilon \sim \mathcal{N}(0, I_d). \quad (33)$$

Comparing (7) and (33), it is not hard to check that

$$Y_\tau \stackrel{(d)}{=} \sqrt{1+\tau} X_{\frac{1}{2} \log(1+\tau)}.$$

Denote $\theta = \frac{1}{2} \log(1+\tau)$ as a short-hand. We have

$$p_{Y_\tau}(x) = p_{\sqrt{1+\tau} X_\theta}(x) \propto p_{X_\theta}\left(\frac{1}{\sqrt{1+\tau}}x\right).$$

Therefore it follows that

$$\nabla \log p_{Y_\tau}(x) = \nabla_x \log p_{X_\theta}\left(\frac{1}{\sqrt{1+\tau}}x\right) = \frac{1}{\sqrt{1+\tau}} s^{\text{cont}}\left(\theta, \frac{1}{\sqrt{1+\tau}}x\right),$$

where we used the definition $s^{\text{cont}}(\theta, \cdot) = \nabla \log p_{X_\theta}(\cdot)$. Plugging the definition $\theta = \frac{1}{2} \log(1+\tau)$ into the above equation yields the desired result. \square

A.2 Proof of Lemma 2

Proof. We first compute the probability density function of z . Recall that $z = w - x_{\text{noisy}}$, thus applying Bayes rule yields

$$p_z(x) = p_w(x + x_{\text{noisy}}) = p^*(x^* = x + x_{\text{noisy}} | x^* + \xi = x_{\text{noisy}}) = \frac{p^*(x + x_{\text{noisy}}) p_\xi(-x)}{p_{x^* + \xi}(x_{\text{noisy}})} \propto p^*(x + x_{\text{noisy}}) p_\xi(-x),$$

where $\xi \sim \mathcal{N}(0, \eta^2 I_d)$. It is straightforward to compute

$$p_\xi(-x) = \frac{1}{(2\pi)^{d/2} \eta^d} e^{-\frac{1}{2\eta^2} \|x\|^2},$$

therefore

$$p_z(x) \propto p^*(x + x_{\text{noisy}}) e^{-\frac{1}{2\eta^2} \|x\|^2}. \quad (34)$$

We proceed to compute the probability density function of Z_τ . According to (24), it follows that

$$\begin{aligned} p_{Z_\tau}(x) &= p_{e^{-\tau} z} * p_{\sqrt{1-e^{-2\tau}} \varepsilon}(x) \\ &= \int p_{e^{-\tau} z}(x') p_{\sqrt{1-e^{-2\tau}} \varepsilon}(x - x') dx' \\ &\propto \int p_z(e^\tau x') \exp\left(-\frac{1}{2(1-e^{-2\tau})} \|x - x'\|^2\right) dx' \\ &\propto \int p^*(x_{\text{noisy}} + e^\tau x') \exp\left(-\frac{1}{2\eta^2} \|e^\tau x'\|^2\right) \exp\left(-\frac{1}{2(1-e^{-2\tau})} \|x - x'\|^2\right) dx', \\ &\propto \int p^*(x') \exp\left(-\frac{1}{2\eta^2} \|x' - x_{\text{noisy}}\|^2\right) \exp\left(-\frac{1}{2(1-e^{-2\tau})} \|x - e^{-\tau}(x' - x_{\text{noisy}})\|^2\right) dx', \end{aligned} \quad (35)$$

where $*$ denotes convolution, the penultimate line follows from (34) and the last line follow from the change of variable $x' \mapsto e^{-\tau}(x' - x_{\text{noisy}})$. One may exercise some brute force to verify that

$$\begin{aligned} &\exp\left(-\frac{1}{2\eta^2} \|x' - x_{\text{noisy}}\|^2\right) \exp\left(-\frac{1}{2(1-e^{-2\tau})} \|x - e^{-\tau}(x' - x_{\text{noisy}})\|^2\right) \\ &= \exp\left(-\frac{e^{2\tau} \|x\|^2}{2(\eta^2 + e^{2\tau} - 1)}\right) \exp\left(-\frac{1}{2(1-e^{-2\tau})} \left\|e^{-\tilde{\tau}} x_{\text{noisy}} + \frac{e^{\tau-\tilde{\tau}} \eta^2 x}{\eta^2 + e^{2\tau} - 1} - e^{-\tilde{\tau}} x'\right\|^2\right) \\ &\propto \exp\left(-\frac{e^{2\tau} \|x\|^2}{2(\eta^2 + e^{2\tau} - 1)}\right) p_{\sqrt{1-e^{-2\tilde{\tau}}} \varepsilon}\left(e^{-\tilde{\tau}} x_{\text{noisy}} + \frac{e^{\tau-\tilde{\tau}} \eta^2 x}{\eta^2 + e^{2\tau} - 1} - e^{-\tilde{\tau}} x'\right), \end{aligned}$$

where $\tilde{\tau}$ is as defined in (27). Plug this back into (35), we see

$$\begin{aligned}
p_{Z_\tau}(x) &\propto \exp\left(-\frac{e^{2\tau}\|x\|^2}{2(\eta^2 + e^{2\tau} - 1)}\right) \int p^*(x') p_{\sqrt{1-e^{-2\tilde{\tau}}}\varepsilon} \left(e^{-\tilde{\tau}}x_{\text{noisy}} + \frac{e^{\tau-\tilde{\tau}}\eta^2 x}{\eta^2 + e^{2\tau} - 1} - e^{-\tilde{\tau}}x'\right) dx' \\
&\propto \exp\left(-\frac{e^{2\tau}\|x\|^2}{2(\eta^2 + e^{2\tau} - 1)}\right) \int p^*(e^{\tilde{\tau}}x') p_{\sqrt{1-e^{-2\tilde{\tau}}}\varepsilon} \left(e^{-\tilde{\tau}}x_{\text{noisy}} + \frac{e^{\tau-\tilde{\tau}}\eta^2 x}{\eta^2 + e^{2\tau} - 1} - x'\right) dx' \\
&\propto \exp\left(-\frac{e^{2\tau}\|x\|^2}{2(\eta^2 + e^{2\tau} - 1)}\right) p_{e^{-\tau}x_0} * p_{\sqrt{1-e^{-2\tilde{\tau}}}\varepsilon} \left(e^{-\tilde{\tau}}x_{\text{noisy}} + \frac{e^{\tau-\tilde{\tau}}\eta^2 x}{\eta^2 + e^{2\tau} - 1}\right) \\
&\propto \exp\left(-\frac{e^{2\tau}\|x\|^2}{2(\eta^2 + e^{2\tau} - 1)}\right) p_{X_{\tilde{\tau}}} \left(e^{-\tilde{\tau}}x_{\text{noisy}} + \frac{e^{\tau-\tilde{\tau}}\eta^2 x}{\eta^2 + e^{2\tau} - 1}\right),
\end{aligned}$$

where the second line applies the change of variable $x' \mapsto e^{\tilde{\tau}}x'$ in the integral, the penultimate line follows from $p_{e^{-\tilde{\tau}}x_0}(x') \propto p^*(e^{\tilde{\tau}}x')$ (since $x_0 \sim p^*$), and the last line follows from $X_{\tilde{\tau}} \stackrel{(d)}{=} e^{-\tilde{\tau}}x_0 + \sqrt{1-e^{-2\tilde{\tau}}}\varepsilon$.

Finally, from the above formula, we obtain

$$\begin{aligned}
\nabla \log p_{Z_\tau}(x) &= \nabla_x \left(-\frac{e^{2\tau}\|x\|^2}{2(\eta^2 + e^{2\tau} - 1)}\right) + \nabla_x \log p_{X_{\tilde{\tau}}} \left(e^{-\tilde{\tau}}x_{\text{noisy}} + \frac{e^{\tau-\tilde{\tau}}\eta^2 x}{\eta^2 + e^{2\tau} - 1}\right) \\
&= -\frac{e^{2\tau}x}{\eta^2 + e^{2\tau} - 1} + \frac{e^{\tau-\tilde{\tau}}\eta^2}{\eta^2 + e^{2\tau} - 1} s^{\text{cont}} \left(\tilde{\tau}, e^{-\tilde{\tau}}x_{\text{noisy}} + \frac{e^{\tau-\tilde{\tau}}\eta^2 x}{\eta^2 + e^{2\tau} - 1}\right),
\end{aligned}$$

where we used the definition $s^{\text{cont}}(\tilde{\tau}, \cdot) = \nabla \log p_{X_{\tilde{\tau}}}(\cdot)$. \square

B Discretization with the exponential integrator

B.1 General form of the exponential integrator

Consider a SDE of the form:

$$dM_\tau = (v(\tau)M_\tau + f(\tau, M_\tau))d\tau + \sqrt{\beta}dB_\tau, \quad \tau \in [0, \tau_\infty], \quad M_0 \sim p_{M_0},$$

where $v : [0, \tau_\infty] \rightarrow \mathbb{R}$, $f : [0, \tau_\infty] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ are deterministic functions, and $\beta > 0$ is a constant. Given discretization time points $0 = \tau_0 \leq \tau_1 \leq \dots \leq \tau_k \leq \tau_\infty$, a naïve way to discretize the SDE is

$$M_{\tau_{i+1}} - M_{\tau_i} \approx (v(\tau_i)M_{\tau_i} + f(\tau_i, M_{\tau_i}))(\tau_{i+1} - \tau_i) + \sqrt{\beta}\sqrt{\tau_{i+1} - \tau_i}\varepsilon_i, \quad i = 0, 1, \dots, k-1,$$

where $\varepsilon_i \sim \mathcal{N}(0, I_d)$ is a standard d -dimensional Gaussian random vector which is independent of M_{τ_i} . Although this approach is straightforward, it has the drawback that the linear term $v(\tau)M_\tau$ is discretized rather crude. For example, for the OU process where $v \equiv -1$, $f \equiv 0$, $\beta = 2$, the SDE can be solved analytically as in (7), while the above approach still has a discretization error.

A more accurate discretization, known to significantly improve the quality of score-based generative models, is given by the *exponential integrator* (Zhang and Chen, 2022), which preserves the linear term and discretizes the SDE to

$$d\widehat{M}_\tau = (v(\tau)\widehat{M}_\tau + f(\tau_i, \widehat{M}_{\tau_i}))d\tau + \sqrt{\beta}dB_\tau, \quad \tau \in [\tau_i, \tau_{i+1}], \quad i = 0, 1, \dots, k,$$

with initialization $\widehat{M}_0 \sim p_{M_0}$. On each time interval $[\tau_i, \tau_{i+1}]$, this is simply a linear SDE, which can be explicitly solved by

$$\widehat{M}_\tau \stackrel{(d)}{=} e^{V(\tau)-V(\tau_i)}\widehat{M}_{\tau_i} + \left(\int_{\tau_i}^\tau e^{V(\tau)-V(\tilde{\tau})}d\tilde{\tau}\right)f(\tau_i, \widehat{M}_{\tau_i}) + \sqrt{\beta}\left(\int_{\tau_i}^\tau e^{2(V(\tau)-V(\tilde{\tau}))}d\tilde{\tau}\right)^{1/2}\varepsilon_i,$$

where V is the antiderivative of v :

$$V(\tau) = \int_0^\tau v(\tilde{\tau})d\tilde{\tau}.$$

Taking $\tau = \tau_{i+1}$, we obtain

$$\begin{aligned}\widehat{M}_{\tau_{i+1}} &\stackrel{(d)}{=} e^{V(\tau_{i+1})-V(\tau_i)} \widehat{M}_{\tau_i} + \left(\int_{\tau_i}^{\tau_{i+1}} e^{V(\tau_{i+1})-V(\tilde{\tau})} d\tilde{\tau} \right) f(\tau_i, \widehat{M}_{\tau_i}) \\ &\quad + \sqrt{\beta} e^{V(\tau_{i+1})} \left(\int_{\tau_i}^{\tau_{i+1}} e^{2(V(\tau_{i+1})-V(\tilde{\tau}))} d\tilde{\tau} \right)^{1/2} \varepsilon_i,\end{aligned}\tag{36}$$

which provides an iterative formula to compute $\widehat{M}_{\tau_{i+1}}$.

B.2 Discretization of DDS-DDPM

Plug the expression of $\nabla \log p_{Y_\tau}$ in Lemma 1 into (21), and use the notation $\tau^{\text{rev}} = \eta^2 - \tau$, we obtain, for $\tau \in [0, \eta^2]$, that

$$\begin{aligned}dY_{\tau^{\text{rev}}}^{\text{rev}} &= \frac{1}{\sqrt{1+\tau^{\text{rev}}}} s^{\text{cont}} \left(\frac{1}{2} \log(1+\tau^{\text{rev}}), \frac{Y_{\tau^{\text{rev}}}^{\text{rev}}}{\sqrt{1+\tau^{\text{rev}}}} \right) d\tau + d\tilde{B}_\tau \\ &= -\frac{1}{\sqrt{\tau^{\text{rev}}}} \varepsilon^{\text{cont}} \left(\frac{1}{2} \log(1+\tau^{\text{rev}}), \frac{Y_{\tau^{\text{rev}}}^{\text{rev}}}{\sqrt{1+\tau^{\text{rev}}}} \right) d\tau + d\tilde{B}_\tau.\end{aligned}$$

Choosing discretization time points. To discretize this SDE, we first choose the discretization time points. Recalling (9), it is most reasonable to discretize at those time points $0 \leq \tau_0^{\text{rev}} \leq \dots \leq \tau_{T'}^{\text{rev}} \leq \eta^2$ which satisfy

$$\frac{1}{2} \log(1+\tau_t^{\text{rev}}) = \frac{1}{2} \log \frac{1}{\bar{\alpha}_t}, \quad 0 \leq t \leq T'.$$

This solves to

$$\tau_t^{\text{rev}} = \bar{\alpha}_t^{-1} - 1.\tag{37}$$

The requirement that $\tau_t^{\text{rev}} \leq \eta^2$ translates to $\bar{\alpha}_t \geq \frac{1}{1+\eta^2}$, which yields the following choice of T' :

$$T' := \max \left\{ t : 0 \leq t \leq T, \bar{\alpha}_t > \frac{1}{\eta^2 + 1} \right\}.\tag{38}$$

Applying the exponential integrator. Now we apply the exponential integrator to discretize the SDE on each time interval $\tau^{\text{rev}} \in [\tau_{t-1}, \tau_t]$, $t = 1, \dots, T'$ as follows:

$$\begin{aligned}d\widehat{Y}_{\tau^{\text{rev}}}^{\text{rev}} &= -\frac{1}{\sqrt{\tau^{\text{rev}}}} \varepsilon^{\text{cont}} \left(\frac{1}{2} \log(1+\tau_t^{\text{rev}}), \frac{\widehat{Y}_{\tau_t^{\text{rev}}}^{\text{rev}}}{\sqrt{1+\tau_t^{\text{rev}}}} \right) d\tau + d\tilde{B}_\tau, \\ &= -\frac{1}{\sqrt{\tau^{\text{rev}}}} \varepsilon_t^* \left(\frac{\widehat{Y}_{\tau_t^{\text{rev}}}^{\text{rev}}}{\sqrt{1+\tau_t^{\text{rev}}}} \right) d\tau + d\tilde{B}_\tau \\ &= -\frac{1}{\sqrt{\tau^{\text{rev}}}} \varepsilon_t^* \left(\sqrt{\bar{\alpha}_t} \widehat{Y}_{\tau_t^{\text{rev}}}^{\text{rev}} \right) d\tau + d\tilde{B}_\tau.\end{aligned}$$

The SDE can be integrated directly on $\tau^{\text{rev}} \in [\tau_{t-1}, \tau_t]$ (see also (36), with $v \equiv 0$), yielding

$$\begin{aligned}\widehat{Y}_{\tau_{t-1}^{\text{rev}}}^{\text{rev}} &= \widehat{Y}_{\tau_t^{\text{rev}}}^{\text{rev}} - 2(\sqrt{\tau_t^{\text{rev}}} - \sqrt{\tau_{t-1}^{\text{rev}}}) \cdot \varepsilon_t^* \left(\sqrt{\bar{\alpha}_t} \widehat{Y}_{\tau_t^{\text{rev}}}^{\text{rev}} \right) + \int_{\eta^2 - \tau_t}^{\eta^2 - \tau_{t-1}} d\tilde{B}_\tau d\tau \\ &\stackrel{(d)}{=} \widehat{Y}_{\tau_t^{\text{rev}}}^{\text{rev}} - 2(\sqrt{\tau_t^{\text{rev}}} - \sqrt{\tau_{t-1}^{\text{rev}}}) \cdot \varepsilon_t^* \left(\sqrt{\bar{\alpha}_t} \widehat{Y}_{\tau_t^{\text{rev}}}^{\text{rev}} \right) + \sqrt{\tau_t^{\text{rev}} - \tau_{t-1}^{\text{rev}}} w_t,\end{aligned}$$

where $w_t \sim \mathcal{N}(0, I_d)$ is independent of $\widehat{Y}_{\tau_t^{\text{rev}}}^{\text{rev}}$. Set $\hat{x}_t = \widehat{Y}_{\tau_t^{\text{rev}}}^{\text{rev}}$, we obtain

$$\hat{x}_{t-1} \stackrel{(d)}{=} \hat{x}_t - 2(\sqrt{\tau_t} - \sqrt{\tau_{t-1}}) \cdot \varepsilon_t^* \left(\sqrt{\bar{\alpha}_t} \hat{x}_t \right) + \sqrt{\tau_t - \tau_{t-1}} w_t, \quad w_t \sim \mathcal{N}(0, I_d),\tag{39}$$

which is exactly the update equation in Algorithm 1, except that ε_t^* is replaced by the noise estimate $\widehat{\varepsilon}_t$.

B.3 Discretization of DDS-DDIM

Plug in the expression of s_Z in Lemma 2 into the probability flow ODE (25), we obtain

$$\begin{aligned} dZ_{\tau^{\text{rev}}}^{\text{rev}} &= \frac{\eta^2 - 1}{\eta^2 + e^{2\tau^{\text{rev}}} - 1} Z_{\tau^{\text{rev}}}^{\text{rev}} d\tau + \frac{e^{\tau^{\text{rev}} - \tilde{\tau}(\tau^{\text{rev}})} \eta^2}{\eta^2 + e^{2\tau^{\text{rev}}} - 1} s^{\text{cont}} \left(\tilde{\tau}(\tau^{\text{rev}}), e^{-\tilde{\tau}(\tau^{\text{rev}})} x_{\text{noisy}} + \frac{e^{\tau^{\text{rev}} - \tilde{\tau}(\tau^{\text{rev}})} \eta^2 x}{\eta^2 + e^{2\tau^{\text{rev}}} - 1} \right) d\tau \\ &= \frac{\eta^2 - 1}{\eta^2 + e^{2\tau^{\text{rev}}} - 1} Z_{\tau^{\text{rev}}}^{\text{rev}} d\tau - \frac{e^{2\tau^{\text{rev}}}}{e^{2\tau^{\text{rev}}} - 1} \varepsilon^{\text{cont}} \left(\tilde{\tau}(\tau^{\text{rev}}), e^{-\tilde{\tau}(\tau^{\text{rev}})} x_{\text{noisy}} + \frac{e^{\tau^{\text{rev}} - \tilde{\tau}(\tau^{\text{rev}})} \eta^2 x}{\eta^2 + e^{2\tau^{\text{rev}}} - 1} \right) d\tau, \end{aligned}$$

where the second line used the definition (27).

Choosing discretization time points. Similar to the derivation in Appendix B.2, we discretize at time points $0 = \tau_0^{\text{rev}} \leq \tau_1^{\text{rev}} \leq \dots \leq \tau_{T'}^{\text{rev}} \leq \eta^2$, which obey

$$\tilde{\tau}(\tau_t^{\text{rev}}) = \frac{1}{2} \log \frac{1}{\bar{\alpha}_t}, \quad t = 0, 1, \dots, T', \quad (40)$$

which solves to

$$\tau_t^{\text{rev}} = \frac{1}{2} \log \frac{\eta^2 + \bar{\alpha}_t - 1}{(\eta^2 + 1)\bar{\alpha}_t - 1}. \quad (41)$$

To make this well-defined, we require

$$\frac{\eta^2 + \bar{\alpha}_t - 1}{(\eta^2 + 1)\bar{\alpha}_t - 1} > 0,$$

which is equivalent to

$$\bar{\alpha}_t > \frac{1}{1 + \eta^2}.$$

This leads to the same choice of T' as in (38). We also set

$$\tau_\infty = \tau_{T'}^{\text{rev}}.$$

It is convenient to introduce a notation for the corresponding discrete schedule of τ_t^{rev} , denoted by

$$\bar{u}_t = e^{-2\tau_t^{\text{rev}}} = \frac{(\eta^2 + 1)\bar{\alpha}_t - 1}{\eta^2 + \bar{\alpha}_t - 1}, \quad t = 0, 1, \dots, T'.$$

Applying the exponential integrator. Now we apply the exponential integrator, which discretizes the ODE on each time interval $\tau^{\text{rev}} \in [\tau_{t-1}, \tau_t]$, $t = 1, \dots, T'$, as

$$\begin{aligned} d\hat{Z}_{\tau^{\text{rev}}}^{\text{rev}} &= \frac{\eta^2 - 1}{\eta^2 + e^{2\tau^{\text{rev}}} - 1} \hat{Z}_{\tau^{\text{rev}}}^{\text{rev}} d\tau - \frac{e^{2\tau^{\text{rev}}}}{e^{2\tau^{\text{rev}}} - 1} \varepsilon^{\text{cont}} \left(\tilde{\tau}(\tau_t^{\text{rev}}), e^{-\tilde{\tau}(\tau_t^{\text{rev}})} x_{\text{noisy}} + \frac{e^{\tau_t^{\text{rev}} - \tilde{\tau}(\tau_t^{\text{rev}})} \eta^2 \hat{Z}_{\tau^{\text{rev}}}^{\text{rev}}}{\eta^2 + e^{2\tau_t^{\text{rev}}} - 1} \right) d\tau \\ &= \frac{\eta^2 - 1}{\eta^2 + e^{2\tau^{\text{rev}}} - 1} \hat{Z}_{\tau^{\text{rev}}}^{\text{rev}} d\tau - \frac{e^{2\tau^{\text{rev}}}}{e^{2\tau^{\text{rev}}} - 1} \varepsilon_t^* \left(\sqrt{\bar{\alpha}_t} x_{\text{noisy}} + \frac{e^{\tau_t^{\text{rev}}} \sqrt{\bar{\alpha}_t} \eta^2 \hat{Z}_{\tau^{\text{rev}}}^{\text{rev}}}{\eta^2 + e^{2\tau_t^{\text{rev}}} - 1} \right) d\tau, \\ &= \frac{\eta^2 - 1}{\eta^2 + e^{2\tau^{\text{rev}}} - 1} \hat{Z}_{\tau^{\text{rev}}}^{\text{rev}} d\tau - \frac{e^{2\tau^{\text{rev}}}}{e^{2\tau^{\text{rev}}} - 1} \varepsilon_t^* \left(\sqrt{\bar{\alpha}_t} x_{\text{noisy}} + \frac{\sqrt{\bar{u}_t} \sqrt{\bar{\alpha}_t} \eta^2 \hat{Z}_{\tau^{\text{rev}}}^{\text{rev}}}{(\eta^2 - 1)\bar{u}_t + 1} \right) d\tau, \end{aligned}$$

where the second line follows from (40), and the last line follows from dividing both the denominator and the numerator in the fraction inside ε_t^* by $e^{2\tau_t^{\text{rev}}}$. This is a first-order linear ODE on $\tau^{\text{rev}} \in [\tau_{t-1}, \tau_t]$, which can be solved explicitly (cf. (36)) by

$$\hat{Z}_{\tau^{\text{rev}}}^{\text{rev}} = \frac{\sqrt{(\eta^2 - 1)e^{-2\tau^{\text{rev}}} + 1}}{\sqrt{(\eta^2 - 1)\bar{u}_t + 1}} \hat{Z}_{\tau_t^{\text{rev}}}^{\text{rev}} + \sqrt{(\eta^2 - 1)e^{-2\tau^{\text{rev}}} + 1} \cdot (h(\eta, e^{-2\tau^{\text{rev}}}) - h(\eta, \bar{u}_t)) \cdot \varepsilon_t^* \left(\sqrt{\bar{\alpha}_t} x_{\text{noisy}} + \frac{\sqrt{\bar{u}_t} \sqrt{\bar{\alpha}_t} \eta^2 \hat{Z}_{\tau^{\text{rev}}}^{\text{rev}}}{(\eta^2 - 1)\bar{u}_t + 1} \right),$$

for $\tau^{\text{rev}} \in [\tau_{t-1}, \tau_t]$, where

$$h(\eta, u) := -\arctan \frac{\eta}{\sqrt{u^{-1} - 1}}.$$

Plug in $\tau^{\text{rev}} = \tau_{t-1}$ in the above solution, and set $z_t = \hat{Z}_{\tau_t^{\text{rev}}}^{\text{rev}}$, we obtain

$$z_{t-1} = \frac{\sqrt{(\eta^2 - 1)\bar{u}_{t-1} + 1}}{\sqrt{(\eta^2 - 1)\bar{u}_t + 1}} z_t + \sqrt{(\eta^2 - 1)\bar{u}_{t-1} + 1} \cdot (h(\eta, \bar{u}_{t-1}) - h(\eta, \bar{u}_t)) \cdot \varepsilon_t^* \left(\sqrt{\bar{\alpha}_t} x_{\text{noisy}} + \frac{\sqrt{\bar{u}_t} \sqrt{\bar{\alpha}_t} \eta^2 z_t}{(\eta^2 - 1)\bar{u}_t + 1} \right). \quad (42)$$

The initialization, which should ideally be $z_{T'} = \hat{Z}_{\tau_\infty}^{\text{rev}} \sim p_{Z_{\tau_\infty}}$, is approximated by $z_{T'} \sim \mathcal{N}(0, I_d)$. This is exactly the update equation and the initialization in Algorithm 2, except that ε_t^* is replaced by the noise estimate $\hat{\varepsilon}_t$.

B.4 Discretization of PCS

We first note that the Metropolis-adjustment step in PCS (cf. Algorithm 4) is standard following the classical form of MALA (Roberts and Rosenthal, 1998). Therefore, we focus on explaining the Langevin step. Recall the continuous-time Langevin dynamics for sampling from the distribution $\exp(\mathcal{L}(\cdot; y) - \frac{1}{2\eta^2} \|\cdot - x\|^2)$:

$$dZ_\tau = -\nabla_{Z_\tau} \mathcal{L}(Z_\tau; y) d\tau + \frac{1}{\eta^2} (Z_\tau - x) d\tau + \sqrt{2} dB_\tau, \quad \tau \geq 0, \quad Z_0 \sim \mathcal{N}(0, I_d). \quad (43)$$

The classical form of MALA, as in Roberts and Rosenthal (1998), performs one step of a straightforward discretization of (43) as the Langevin step, as follows:

$$z_{n+\frac{1}{2}} \approx z_n - \gamma \nabla_{z_n} \mathcal{L}(z_n; y) + \frac{\gamma}{\eta^2} (z_n - x) + \sqrt{2\gamma} w_n, \quad w_n \sim \mathcal{N}(0, I_d).$$

In our setting, due to the presence of the linear drift term $\frac{1}{\eta^2} (Z_\tau - x)$, which can be quite large when η is small, we apply the exponential integrator instead. Set the discretization time points $\tau_n = n\gamma$, the exponential integrator reads as

$$dZ_\tau = -\nabla_{Z_{n\gamma}} \mathcal{L}(Z_{n\gamma}; y) d\tau + \frac{1}{\eta^2} (Z_\tau - x) d\tau + \sqrt{2} dB_\tau, \quad n\gamma \leq \tau \leq (n+1)\gamma.$$

Solve this linear SDE on $n\gamma \leq \tau \leq (n+1)\gamma$ directly (see also (36)) to obtain

$$Z_{(n+1)\gamma} \stackrel{(d)}{=} r Z_{n\gamma} + (1-r)x + \eta^2(1-r) \nabla_{Z_{n\gamma}} \mathcal{L}(Z_{n\gamma}; y) + \eta \sqrt{1-r^2} w_n, \quad w_n \sim \mathcal{N}(0, I_d),$$

where $r := e^{-\gamma/\eta^2}$. This is the same as the update equation for the Langevin step in PCS (cf. Algorithm 4).

C Proof of main theorems

C.1 Proof of Theorem 1

Proof. The proof is based on two lemmas on the one-step transition kernel of DPnP and the asymptotic behavior of the transition kernel, which we will present soon. First, we set up some notations. Denote

$$p_\eta(x) := p_{x^* \sim p^*, \varepsilon \sim \mathcal{N}(0, I_d)}(x^* + \eta\varepsilon = x) = \frac{1}{(2\pi)^{d/2} \eta^d} \int p^*(z) e^{-\frac{1}{2\eta^2} \|x-z\|^2} dz.$$

From the first equality, it is clear that $p_\eta \rightarrow p^*$ when $\eta \rightarrow 0^+$. We will also use the notation q_η defined in (31), which we recall here:

$$q_\eta(x) := \frac{1}{(2\pi)^{d/2} \eta^d} \int e^{\mathcal{L}(z; y) - \frac{1}{2\eta^2} \|x-z\|^2} dz.$$

In virtue of the Assumption 1, we know that q_η is finite for all $x \in \mathbb{R}^d$.

For convenience, we introduce a notation for application of transition kernels. For a probability distribution $p(x)$ and a probability transition kernel $K(x, x')$, denote by $p \circ K$ the probability distribution given by

$$p \circ K(x') = \int p(x)K(x, x')dx.$$

The first lemma characterizes the one-step behavior of DPnP in terms of Markov transition kernels.

Lemma 5. *Under the settings of Lemma 4 and Lemma 3, the one-step transition kernel of DPnP with $\eta_k = \eta$ is given by:*

$$K_{\text{DPnP}, \eta}(x, x') = \left(\int \frac{q_0(z)}{p_\eta(z)} e^{-\frac{1}{2\eta^2} \|z-x\|^2 - \frac{1}{2\eta^2} \|z-x'\|^2} dz \right) \frac{p^*(x')}{q_\eta(x)}.$$

In other words, if \hat{x}_k has distribution $p_{\hat{x}_k}$, then the distribution of \hat{x}_{k+1} is

$$p_{\hat{x}_{k+1}}(x') = p_{\hat{x}_k} \circ K_{\text{DPnP}, \eta}(x) = \int p_{\hat{x}_k}(x) K_{\text{DPnP}, \eta}(x, x') dx.$$

The proof is postponed to Appendix C.3. The next lemma analyzes the ergodic properties of the Markov chain with transition kernel $K_{\text{DPnP}, \eta}$. These properties are known (Bouman and Buzzard, 2023) but scattered in different literatures, so we will provide a brief proof to be self-contained.

Lemma 6. *The Markov transition kernel $K_{\text{DPnP}, \eta}$ has the following properties:*

(i) (Stationary distribution.) Let π_η be the probability distribution defined by

$$\pi_\eta(x) = c_\eta p^*(x) q_\eta(x),$$

where $c_\eta > 0$ is the normalization constant such that $\int \pi_\eta(x) dx = 1$. Then $K_{\text{DPnP}, \eta}$ is reversible with stationary distribution π_η .

(ii) (Convergence.) For any initial distribution p , the distribution of the Markov chain with kernel $K_{\text{DPnP}, \eta}$ converges to π_η :

$$\text{TV}(p \circ K_{\text{DPnP}, \eta}^{(n)}, \pi_\eta) \rightarrow 0, \quad n \rightarrow \infty, \quad (44)$$

where $K_{\text{DPnP}, \eta}^{(n)}$ is the n -step transition kernel of $K_{\text{DPnP}, \eta}$.

The proof is postponed to Appendix C.4. We now show how to prove Theorem 1 with the above two lemmas. With the annealing schedule in Theorem 1, between steps $k_{l-1} \leq k < k_l$, which consist of consecutive $(k_l - k_{l-1})$ steps, the transition kernel of one-step of DPnP is $K_{\text{DPnP}, \varepsilon_l}$. As $(k_l - k_{l-1}) \rightarrow \infty$, Lemma 6 implies that

$$\text{TV}(p_{\hat{x}_{k_l}}, \pi_{\varepsilon_l}) = \text{TV}(p_{\hat{x}_{k_{l-1}}} \circ K_{\text{DPnP}, \varepsilon_l}^{(k_l - k_{l-1})}, \pi_{\varepsilon_l}) \rightarrow 0.$$

Under the assumption in Theorem 1 that $\varepsilon_l \rightarrow 0$, we let $l \rightarrow \infty$ to see $\lim_{l \rightarrow \infty} \pi_{\varepsilon_l} = c_0 p^*(\cdot) e^{\mathcal{L}(\cdot; y)} = p^*(\cdot | y)$, thus $p_{\hat{x}_{k_l}} \rightarrow p^*(\cdot | y)$, as claimed. \square

C.2 Proof of Lemma 4

Proof. For DDS-DDPM, we note that under the continuous-time limit in Lemma 4, the discretization time points given by (37) verify

$$\tau_0^{\text{rev}} = 0, \quad \sup_{0 \leq t \leq T'-1} |\tau_t^{\text{rev}} - \tau_{t+1}^{\text{rev}}| \rightarrow 0, \quad \tau_{T'}^{\text{rev}} \rightarrow \left(\frac{1}{1 + \eta^2} \right)^{-1} - 1 = \eta^2, \quad T' \rightarrow \infty.$$

Therefore, these discretization time points $0 = \tau_0^{\text{rev}} \leq \dots \leq \tau_{T'}^{\text{rev}} \leq \eta^2$ form a partition of $[0, \eta^2]$, which becomes infinitely fine in the continuous-time limit. Thus the discretized integrator (39) converges to the solution of the SDE (21), which, as we have already argued in Appendix A, produces samples obeying the denoising posterior distribution $p^*(\cdot | x_{\text{noisy}})$, as claimed.

The proof for DDS-DDPM follows similarly, by observing that the discretization time points in (37) form an infinitely fine partition of $[0, \infty)$ in the continuous-time limit. \square

C.3 Proof of Lemma 5

Proof. The proof is based on computing the transition kernel of the two subroutines. We claim that

- (i) Sampling with probability density proportional to $\exp(\mathcal{L}(\cdot; y) - \frac{1}{2\eta^2} \|\cdot - x\|^2)$ is equivalent to applying the following Markov transition kernel

$$K_{\text{PCS},\eta}(x, x') = \frac{1}{q_\eta(x)} e^{\mathcal{L}(x'; y) - \frac{1}{2\eta^2} \|x' - x\|^2}.$$

- (ii) Sampling with probability $p^*(x^* | x^* + \eta\varepsilon = x)$, where $\varepsilon \sim \mathcal{N}(0, I_d)$, is equivalent to applying the following Markov transition kernel:

$$K_{\text{DDS},\eta}(x, x') = \frac{1}{p_\eta(x)} p^*(x') e^{-\frac{1}{2\eta^2} \|x' - x\|^2}.$$

It is then clear that

$$K_{\text{DPnP},\eta}(x, x') = \int K_{\text{PCS},\eta}(x, z) K_{\text{DDS},\eta}(z, x') dz = \left(\int \frac{q_0(z)}{p_\eta(z)} e^{-\frac{1}{2\eta^2} \|z - x\|^2 - \frac{1}{2\eta^2} \|z - x'\|^2} dz \right) \frac{p^*(x')}{q_\eta(x)},$$

as desired. We now prove the above two claims. For (i), note that by (30), we know $K_{\text{DDS},\eta}(x, \cdot) \propto p^*(\cdot) e^{-\frac{1}{2\eta^2} \|\cdot - x\|^2}$. Thus it suffices to compute the normalization constant, which is

$$\int p^*(x') e^{-\frac{1}{2\eta^2} \|x' - x\|^2} dx' = p_\eta(x),$$

by the definition of p_η . Therefore

$$K_{\text{DDS},\eta}(x, x') = \frac{1}{p_\eta(x)} p^*(x') e^{-\frac{1}{2\eta^2} \|x' - x\|^2},$$

as claimed. The proof of (ii) follows similarly. \square

C.4 Proof of Lemma 6

Proof. We first introduce a fundamental lemma (Polyanskiy and Wu, 2024), which provides a simple method to bound the total variation between two distributions.

Lemma 7 (Data-processing inequality). *Let p, q be two probability distributions, and K be a probability transition kernel. Then*

$$\text{TV}(p \circ K, q \circ K) \leq \text{TV}(p, q).$$

We now prove the two items in Lemma 6 separately.

Proof of (i). We first show that π_η is well-defined, i.e., $\int p^*(x) q_\eta(x) dx < \infty$. This can be seen from Assumption 1, which implies $q_\eta(x) \lesssim \int e^{-\frac{1}{2\eta^2} \|x - z\|^2} dz \lesssim 1$, hence

$$\int p^*(x) q_\eta(x) dx \lesssim \int p^*(x) dx = 1.$$

To show that $K_{\text{DPnP},\eta}$ is reversible with stationary distribution π_η , it suffices to verify

$$\pi_\eta(x) K_{\text{DPnP},\eta}(x, x') = \pi_\eta(x') K_{\text{DPnP},\eta}(x', x), \quad \forall x, x' \in \mathbb{R}^d.$$

However, it is easily checked that both sides are equal to

$$c_\eta \left(\int \frac{q_0(z)}{p_\eta(z)} e^{-\frac{1}{2\eta^2} \|z - x\|^2 - \frac{1}{2\eta^2} \|z - x'\|^2} dz \right) p^*(x') p^*(x).$$

Proof of (ii). We define an auxiliary Markov transition kernel $K_{\text{aux},\eta} = K_{\text{DDS},\eta} \circ K_{\text{PCS},\eta}$. More explicitly,

$$K_{\text{aux},\eta}(x, x') = \int K_{\text{DDS},\eta}(x, z) K_{\text{PCS},\eta}(z, x') dz = \left(\int \frac{p^*(z)}{q_\eta(z)} e^{-\frac{1}{2\eta^2}\|z-x\|^2 - \frac{1}{2\eta^2}\|z-x'\|^2} dz \right) \frac{e^{\mathcal{L}(x';y)}}{p_\eta(x)}. \quad (45)$$

It is easy to see that

$$p \circ K_{\text{DPnP},\eta}^{(n)} = p \circ K_{\text{PCS},\eta} \circ K_{\text{aux},\eta}^{(n-1)} \circ K_{\text{DDS},\eta}. \quad (46)$$

Thus we are led to investigate the ergodic properties of $K_{\text{aux},\eta}$. Similar to the proof of item (i) above, it is not hard to show that $K_{\text{aux},\eta}$ is reversible with respect to the stationary distribution

$$\mu_\eta(x) := c_\eta p_\eta(x) q_0(x) = c_\eta p_\eta(x) e^{\mathcal{L}(x;y)}.$$

Moreover, one may check that

$$\pi_\eta = \mu_\eta \circ K_{\text{DDS},\eta}. \quad (47)$$

It is apparent that $\mu(x') > 0$ and $K_{\text{aux},\eta}(x, x')/\mu_\eta(x') > 0$ for all $x, x' \in \mathbb{R}^d$. By Tierney (1994, Corollary 1), such a Markov transition kernel obeys, for any probability distribution q , that

$$\text{TV}(q \circ K_{\text{aux},\eta}^{(n)}, \mu_\eta) \rightarrow 0, \quad n \rightarrow \infty.$$

In view of (46) and (47), we set $q = p \circ K_{\text{PCS},\eta}$ and invoke the data-processing inequality to obtain

$$\begin{aligned} \text{TV}(p \circ K_{\text{DPnP},\eta}^{(n)}, \pi_\eta) &= \text{TV}(q \circ K_{\text{aux},\eta}^{(n-1)} \circ K_{\text{DDS},\eta}, \mu_\eta \circ K_{\text{DDS},\eta}) \\ &\leq \text{TV}(q \circ K_{\text{aux},\eta}^{(n-1)}, \mu_\eta) \\ &\rightarrow 0, \end{aligned}$$

as $n \rightarrow \infty$. This completes the proof. \square

C.5 Proof of Theorem 2

Proof. Denote by $\tilde{K}_{\text{PCS},\eta}$ and $\tilde{K}_{\text{DDS},\eta}$ and the transition kernels for PCS and for DDS, respectively. Note that these may deviate from the transition kernels $K_{\text{PCS},\eta}$ and $K_{\text{DDS},\eta}$ defined for the idealized asymptotic setting in Appendix C. We have

$$\begin{aligned} \text{TV}(p_{\hat{x}_N}, \pi_\eta) &= \text{TV}(p_{\hat{x}_{N-\frac{1}{2}}} \circ \tilde{K}_{\text{DDS},\eta}, \pi_\eta) \\ &\leq \text{TV}(p_{\hat{x}_{N-\frac{1}{2}}} \circ K_{\text{DDS},\eta}, \pi_\eta) + \text{TV}(p_{\hat{x}_{N-\frac{1}{2}}} \circ K_{\text{DDS},\eta}, p_{\hat{x}_{N-\frac{1}{2}}} \circ \tilde{K}_{\text{DDS},\eta}) \\ &\leq \text{TV}(p_{\hat{x}_{N-\frac{1}{2}}} \circ K_{\text{DDS},\eta}, \pi_\eta) + \varepsilon_{\text{DDS}}, \end{aligned}$$

where the second line is triangle inequality, and the third line follows from the assumption in Theorem 2 that DDS has error at most ε_{DDS} in total variation, by taking the input of DDS to be $\hat{x}_{N-\frac{1}{2}}$.

Similarly, from $p_{\hat{x}_{N-\frac{1}{2}}} = p_{\hat{x}_{N-1}} \circ \tilde{K}_{\text{PCS},\eta}$ and the assumption that PCS has error at most ε_{PCS} in total variation, we can show

$$\text{TV}(p_{\hat{x}_{N-\frac{1}{2}}} \circ K_{\text{DDS},\eta}, \pi_\eta) \leq \text{TV}(p_{\hat{x}_{N-1}} \circ K_{\text{PCS},\eta} \circ K_{\text{DDS},\eta}, \pi_\eta) + \varepsilon_{\text{PCS}} = \text{TV}(p_{\hat{x}_{N-1}} \circ K_{\text{DPnP},\eta}, \pi_\eta) + \varepsilon_{\text{PCS}}.$$

The above two inequalities together imply

$$\text{TV}(p_{\hat{x}_N}, \pi_\eta) \leq \text{TV}(p_{\hat{x}_{N-1}} \circ K_{\text{DPnP},\eta}, \pi_\eta) + \varepsilon_{\text{DDS}} + \varepsilon_{\text{PCS}}.$$

Iterating this process, we obtain

$$\text{TV}(p_{\hat{x}_N}, \pi_\eta) \leq \text{TV}(p_{\hat{x}_1} \circ K_{\text{DPnP},\eta}^{(N-1)}, \pi_\eta) + (N-1)(\varepsilon_{\text{DDS}} + \varepsilon_{\text{PCS}}). \quad (48)$$

It remains to bound $\text{TV}(p_{\hat{x}_1} \circ K_{\text{DPnP},\eta}^{(N-1)}, \pi_\eta)$. For this, we need the following two lemmas.

Lemma 8 (Comparing TV and χ^2 -divergence, [Polyanskiy and Wu \(2024\)](#)). *For any two distributions p, q , we have*

$$\text{TV}(p, q) \leq \sqrt{\chi^2(p \| q)}.$$

Lemma 9 (χ^2 -contractivity of $K_{\text{DPnP}, \eta}$). *There exists some $\lambda := \lambda(p^*, \mathcal{L}, \eta) \in (0, 1)$, such that for any probability distribution $p(x)$, we have*

$$\chi^2(p \circ K_{\text{DPnP}, \eta}^{(N)} \| \pi_\eta) \leq \lambda^{2N} \chi^2(p \| \pi_\eta).$$

A form of Lemma 9 is well-known for Markov chains with countable state spaces, but relatively few sources provide a complete proof for the abstract setting we consider here with continuous state space. For sake of completeness, we prove Lemma 9 in Appendix C.6.

Combining the above two lemmas, we obtain

$$\text{TV}(p_{\hat{x}_1} \circ K_{\text{DPnP}, \eta}^{(N-1)}, \pi_\eta) \leq \sqrt{\chi^2(p_{\hat{x}_1} \circ K_{\text{DPnP}, \eta}^{(N-1)} \| \pi_\eta)} \leq \lambda^{N-1} \sqrt{\chi^2(p_{\hat{x}_1} \| \pi_\eta)}.$$

Plug this into (48), we obtain

$$\text{TV}(p_{\hat{x}_N}, \pi_\eta) \leq \lambda^{N-1} \sqrt{\chi^2(p_{\hat{x}_1} \| \pi_\eta)} + (N-1)(\varepsilon_{\text{DDS}} + \varepsilon_{\text{PCS}}).$$

With $N \asymp \frac{\log(1/\varepsilon_{\text{acc}})}{1-\lambda}$ such that $\lambda^{N-1} \leq \exp(-(N-1)(1-\lambda)) \leq \varepsilon_{\text{acc}}$, the desired result readily follows. \square

C.6 Proof of Lemma 9

Proof. We need a few fundamental properties of reversible Markov chains, which are collected below.

First we set up some notations. Define the Hilbert space $L^2(\pi)$ to be the space of square-integrable functions with respect to measure π , i.e., those functions $f : \mathbb{R}^d \rightarrow \mathbb{C}$ such that

$$\|f\|_{L^2(\pi)} := \left(\int |f(x)|^2 \pi(x) dx \right)^{1/2} < \infty.$$

The first well-known property ([Saloff-Coste, 1997](#)) offers a way to represent a reversible transition kernel as a self-adjoint operator (infinite-dimensional symmetric matrix).

Lemma 10 (Self-adjoint representation of reversible Markov operator). *Assume $K(x, x')$ is a Markov transition kernel that is reversible with respect to the stationary distribution $\pi(x)$. Then the integral operator $\mathcal{K} : L^2(\pi) \rightarrow L^2(\pi)$ defined by*

$$\mathcal{K}f(x) = \int K(x, x') f(x') dx'$$

is self-adjoint and compact. For any probability distribution $p(x)$ such that $\int \frac{p^2(x)}{\pi(x)} dx < \infty$, we have

$$\int p(x) \cdot \mathcal{K}f(x) dx = \int p \circ K(x') f(x') dx'.$$

Moreover, the eigenvalues of \mathcal{K} are the same as those of K .

The following theorem is a generalization of the classical Perron-Frobenius theory for finite-dimensional transition matrix to strictly positive operators. The form we present here can be found in [Schaefer \(2012, Theorem V.6.6\)](#); see also [Bourbaki \(2023b, Theorem III.6.7\)](#) for a more elementary treatment which can also be adapted to the form we need.

Theorem 3 (Jentzsch). *Let $K(x, x')$ be a Markov transition kernel. If $K(x, x') > 0$ for any $x, x' \in \mathbb{R}^d$, then K has a unique stationary distribution π . Moreover, 1 is a simple eigenvalue of K , with π being the only left eigenfunction, and the constant function 1 being the only right eigenfunction. In addition, there exists $\lambda \in (0, 1)$ such that any other eigenvalue of K has modulus no larger than λ .*

We are now ready to prove Lemma 9. We divide the proof into the following steps.

Step 1: controlling the eigenvalues of $\mathcal{K}_{\text{DPnP},\eta}$. Recall the auxiliary kernel $K_{\text{aux},\eta}$ defined in (45). It is a standard result in linear algebra or function analysis (Bourbaki, 2023a) that $K_{\text{aux},\eta} = K_{\text{PCS},\eta} \circ K_{\text{DDS},\eta}$ has same eigenvalues as $K_{\text{DPnP},\eta} = K_{\text{DDS},\eta} \circ K_{\text{PCS},\eta}$. From (45), it is easy to check $K_{\text{aux},\eta}(x, x') > 0$, thus Theorem 3 implies 1 is a simple eigenvalue of $\mathcal{K}_{\text{DPnP},\eta}$. Moreover, there exists $\lambda := \lambda(p^*, \mathcal{L}, \eta) \in (0, 1)$, such that any other eigenvalue of $K_{\text{aux},\eta}$ has modulus no larger than λ .

Since $K_{\text{DPnP},\eta}$ has the same eigenvalues as $K_{\text{aux},\eta}$, and, by Lemma 10, the operator $\mathcal{K}_{\text{DPnP},\eta}$ also has the same eigenvalues as these two, we conclude that $\mathcal{K}_{\text{DPnP},\eta}$ is a self-adjoint compact operator on $L^2(\pi_\eta)$, of whom 1 is a simple eigenvalue. Moreover, any other eigenvalue of $\mathcal{K}_{\text{DPnP},\eta}$ has modulus no larger than λ .

Step 2: establishing the contractivity of $\mathcal{K}_{\text{DPnP},\eta}$ in $L^2(\pi_\eta)$. It is easy to verify that the constant function $\mathbf{1}$, which takes value 1 for any $x \in \mathbb{R}^d$, is a eigenfunction of $\mathcal{K}_{\text{DPnP},\eta}$ associated to the simple eigenvalue 1, thus is the only (up to scaling) eigenfunction associated to that eigenvalue. It is also a unit-length eigenfunction, since $\|\mathbf{1}\|_{L^2(\pi_\eta)} = (\int \mathbf{1} \cdot \pi_\eta(x) dx)^{1/2} = 1$. Therefore, the operator $\mathcal{K}_{\text{DPnP},\eta} - \mathbf{1}\mathbf{1}^\top$ is a self-adjoint operator whose eigenvalues have modulus no larger than λ , where $\mathbf{1}\mathbf{1}^\top$ is the orthogonal projection onto $\mathbf{1}$ in $L^2(\pi_\eta)$, defined by

$$\mathbf{1}\mathbf{1}^\top f(x) \equiv \int f(x') \pi_\eta(x') dx', \quad \forall x \in \mathbb{R}^d.$$

Using the fact that $\mathcal{K}_{\text{DPnP},\eta} \mathbf{1}\mathbf{1}^\top = \mathbf{1}\mathbf{1}^\top \mathcal{K}_{\text{DPnP},\eta} = \mathbf{1}\mathbf{1}^\top$, one may show $(\mathcal{K}_{\text{DPnP},\eta} - \mathbf{1}\mathbf{1}^\top)^N = \mathcal{K}_{\text{DPnP},\eta}^{(N)} - \mathbf{1}\mathbf{1}^\top$ by expanding the product, see e.g. Saloff-Coste (1997). Consequently, $\mathcal{K}_{\text{DPnP},\eta}^N - \mathbf{1}\mathbf{1}^\top$ is a self-adjoint operator whose eigenvalues have modulus no larger than λ^N , i.e.,

$$\|\mathcal{K}_{\text{DPnP},\eta}^N - \mathbf{1}\mathbf{1}^\top\|_{L^2(\pi_\eta) \rightarrow L^2(\pi_\eta)} \leq \lambda^N, \quad (49)$$

where $\|\cdot\|_{L^2(\pi_\eta) \rightarrow L^2(\pi_\eta)}$ denotes the operator norm on $L^2(\pi_\eta)$.

Step 3: bounding the inner product of $p \circ K_{\text{DPnP},\eta}^{(N)} - \pi_\eta$ with any square-integrable function. Note that when $\chi^2(p \parallel \pi_\eta) = \infty$, the conclusion is trivially true. For the rest part of the proof, we assume $\chi^2(p \parallel \pi_\eta) < \infty$. Now, for any $f \in L^2(\pi_\eta)$, by applying Lemma 10 iteratively, we obtain

$$\begin{aligned} \int p \circ K_{\text{DPnP},\eta}^{(N)}(x) f(x) dx &= \int p(x') \mathcal{K}_{\text{DPnP},\eta}^N f(x') dx' \\ &= \int p(x') \cdot (\mathcal{K}_{\text{DPnP},\eta}^N - \mathbf{1}\mathbf{1}^\top) f(x') dx' + \int p(x') \mathbf{1}\mathbf{1}^\top f(x') dx' \\ &= \int p(x') \cdot (\mathcal{K}_{\text{DPnP},\eta}^N - \mathbf{1}\mathbf{1}^\top) f(x') dx' + \int f(x') \pi_\eta(x') dx', \end{aligned}$$

where the last line follows from the definition of $\mathbf{1}\mathbf{1}^\top$ and $\int p(x') dx' = 1$. Rearrange the terms to see

$$\int (p \circ K_{\text{DPnP},\eta}^{(N)}(x) - \pi_\eta(x)) f(x) dx = \int p(x') \cdot (\mathcal{K}_{\text{DPnP},\eta}^N - \mathbf{1}\mathbf{1}^\top) f(x') dx'. \quad (50)$$

In particular, taking $p = \pi_\eta$ yields

$$0 = \int \pi_\eta(x') \cdot (\mathcal{K}_{\text{DPnP},\eta}^N - \mathbf{1}\mathbf{1}^\top) f(x') dx'. \quad (51)$$

Subtract (51) from (50), and then take absolute value, we obtain

$$\begin{aligned} \left| \int (p \circ K_{\text{DPnP},\eta}^{(N)}(x) - \pi_\eta(x)) f(x) dx \right| &= \left| \int (p(x') - \pi_\eta(x')) \cdot (\mathcal{K}_{\text{DPnP},\eta}^N - \mathbf{1}\mathbf{1}^\top) f(x') dx' \right| \\ &\leq \left(\int \frac{(p(x') - \pi_\eta(x'))^2}{\pi_\eta(x)} dx \right)^{1/2} \cdot \|(\mathcal{K}_{\text{DPnP},\eta}^N - \mathbf{1}\mathbf{1}^\top) f(x')\|_{L^2(\pi_\eta)} \\ &\leq \sqrt{\chi^2(p \parallel \pi_\eta)} \cdot \lambda^N \|f\|_{L^2(\pi_\eta)}. \end{aligned} \quad (52)$$

Step 4: choosing an appropriate square-integrable function. Now, set

$$f(x) = \frac{p \circ K_{\text{DPnP},\eta}^{(N)}(x) - \pi_\eta(x)}{\pi_\eta(x)}.$$

It is easily checked that

$$\begin{aligned} \int (p \circ K_{\text{DPnP},\eta}^{(N)}(x) - \pi_\eta(x))f(x)dx &= \chi^2(p \circ K_{\text{DPnP},\eta}^{(N)} \parallel \pi_\eta), \\ \|f\|_{L^2(\pi_\eta)} &= \sqrt{\chi^2(p \circ K_{\text{DPnP},\eta}^{(N)} \parallel \pi_\eta)}. \end{aligned}$$

Plug these equations into (52), we obtain

$$\chi^2(p \circ K_{\text{DPnP},\eta}^{(N)} \parallel \pi_\eta) \leq \lambda^{2N} \chi^2(p \parallel \pi_\eta),$$

as claimed. \square

D Additional experiment details

D.1 Hyperparameters of competing methods

We made our best effort to fine-tune the other algorithms within a reasonable amount of time for each task. We list the parameters, following the notation in the original papers (Chung et al., 2023a; Song et al., 2023b), as follows.

- Phase retrieval: For DPS, the stepsize is set to 0.8. For LGD-MC, the MC sampling variance $r_t = 0.05$, the loss coefficient $\lambda = 10^{-3}$, and the learning rate is set to 400.0.
- Quantized sensing: For DPS, the stepsize is set to 100.0. For LGD-MC, the MC sampling variance $r_t = 0.05$, the loss coefficient $\lambda = 2 \times 10^{-5}$, and the learning rate is set to 500.0.
- Super-resolution: For DPS, the learning rate is set to 0.6. For LGD-MC, the MC sampling variance $r_t = 0.05$, the loss coefficient $\lambda = 10^{-3}$, and the learning rate is set to 60.0.

D.2 Additional performance metrics

Computation time in terms of Neural Function Estimations (NFEs). In addition to the clock time statistics, we also measure the computational cost per sample of different algorithms in terms of the number of Neural Function Estimations, i.e., the number of calls to score functions. The results are in Table 3. Note that the NFEs for DPnP depends on the initialization, the annealing schedule (and the number of timesteps for DPnP-DDIM). We provide typical numbers of NFEs with the choice of parameters used in this paper.

Table 3: Number of NFEs for different algorithms.

Algorithm	DPnP-DDIM	DPnP-DDPM	DPS	LGD-MC
NFEs	~ 1500	~ 3000	1000	1000

Frechet Inception Distance (FID) and Structural Similarity Index Measure (SSIM). We also compare the FID and SSIM of different algorithms across different tasks. The results are shown in Table 4 and Table 5. It should be pointed out that FID is arguably not very relevant to measure the quality of solving inverse problems, as accurately solving inverse problems means that the generated distribution is close to the *conditional*, i.e., *posterior* distribution of the image, while FID only measure the closeness to the *unconditional*, i.e., *prior* distribution of the image.

Table 4: FID and SSIM of solving inverse problems on FFHQ 256×256 validation dataset (1k samples).

Algorithm	Super-resolution (4x, linear)		Phase retrieval (nonlinear)		Quantized sensing (nonlinear)	
	FID ↓	SSIM ↑	FID ↓	SSIM ↑	FID ↓	SSIM ↑
DPnP-DDIM (ours)	36.3	0.668	46.5	0.631	37.3	0.712
DPS (Chung et al., 2023a)	38.6	0.636	52.0	0.494	42.1	0.601
LGD-MC ($n = 5$) (Song et al., 2023b)	36.8	0.651	82.3	0.414	40.3	0.639

Table 5: FID and SSIM of solving inverse problems on ImageNet 256×256 validation dataset (1k samples).

Algorithm	Super-resolution (4x, linear)		Phase retrieval (nonlinear)		Quantized sensing (nonlinear)	
	FID ↓	SSIM ↑	FID ↓	SSIM ↑	FID ↓	SSIM ↑
DPnP-DDIM (ours)	47.5	0.510	73.5	0.289	43.2	0.623
DPS (Chung et al., 2023a)	61.4	0.496	92.7	0.318	82.4	0.459
LGD-MC ($n = 5$) (Song et al., 2023b)	46.2	0.503	89.8	0.234	46.8	0.563