
Solving General Noisy Inverse Problem via Posterior Sampling: A Policy Gradient Viewpoint

Haoyue Tang^{1,*}
Hanyu Wang^{3,†}
¹Meta AI

Tian Xie¹
Chenyang Zhang¹
²Yale University

Aosong Feng²
Yang Bai¹
³University of Maryland, College Park

Abstract

Solving image inverse problems (e.g., super-resolution and inpainting) requires generating a high fidelity image that matches the given input (the low-resolution image or the masked image). By using the input image as guidance, we can leverage a pretrained diffusion generative model to solve a wide range of image inverse tasks without task specific model fine-tuning. To precisely estimate the guidance score function of the input image, we propose Diffusion Policy Gradient (DPG), a tractable computation method by viewing the intermediate noisy images as policies and the target image as the states selected by the policy. Experiments show that our method is robust to both Gaussian and Poisson noise degradation on multiple linear and non-linear inverse tasks, resulting into a higher image restoration quality on FFHQ, ImageNet and LSUN datasets.

1 Introduction and Problem Formulation

Denosing Diffusion Probabilistic Models (DDPM) Ho et al. (2020); Sohl-Dickstein et al. (2015) provide tractable solutions to model an unknown high quality image distribution. Their modeling and generation capabilities have been exploited in a wide range of image inverse problems Dhariwal and Nichol (2021); Blattmann et al. (2022); Rombach et al. (2021); Kawar

[†]Work done during an internship at Meta AI.

*Email: tanghaoyue13@tsinghua.org.cn

et al. (2022), where the objective is to recover a high-quality image corresponding to a given low-resolution or blurred image. However, training a diffusion model from scratch for each inverse task is time-consuming. Alternatively, one can use the input image as a guidance, and recover the high-quality image using guided diffusion using a pretrained diffusion generative model Ho and Salimans (2021); Dhariwal and Nichol (2021). Nonetheless, when the input image is distorted by random noise, this guidance signal becomes inaccurate. Therefore, solving such noisy inverse problems is challenging.

We now describe the noisy image inverse problem in more details. Consider \mathbf{x}_0 as a high-quality image with distribution $p_0(\mathbf{x}_0)$. Let \mathbf{y} be a noisy input image obtained by applying an operator \mathcal{A} to image \mathbf{x}_0 , i.e.,

$$\mathbf{y} = \mathcal{A}(\mathbf{x}_0) + \mathbf{n}, \quad (1)$$

where \mathbf{n} represents distorted random noise. The operator \mathcal{A} is dependent on the specific image inverse tasks. For example, in super-resolution tasks, \mathcal{A} is the downsampling operator; in inpainting problems, the operator \mathcal{A} extracts the unmasked pixels of an image. However, the operator \mathcal{A} is often low-rank or invertible, making direct computation of the inverse \mathbf{y} impossible. Alternatively, one can leverage the Bayes' rule $p_0(\mathbf{x}_0|\mathbf{y}) \propto p_0(\mathbf{x}_0)p_0(\mathbf{y}|\mathbf{x}_0)$ to sample from both the prior $p_0(\mathbf{x}_0)$ and the likelihood $p_0(\mathbf{y}|\mathbf{x}_0)$, where the prior $p_0(\mathbf{x}_0)$ is implicitly modeled by a pre-trained diffusion generative model.

There are currently two lines of work in utilizing pre-trained diffusion generative models to solve image inverse problems. The first line of work utilizes the low rank structure of the operator \mathcal{A} , and directly plugs the known information \mathbf{y} into the corresponding space of \mathbf{x}_0 . SDEdit Meng et al. (2022) solves image inpainting and stroke based generation tasks by plugging a noisy \mathbf{y} into a selected starting point of the diffusion generation process. Blended Diffusion Avrahami et al. (2022, 2023) and DiffEdit Couairon et al. (2023) enhance image inpainting and editing performance by

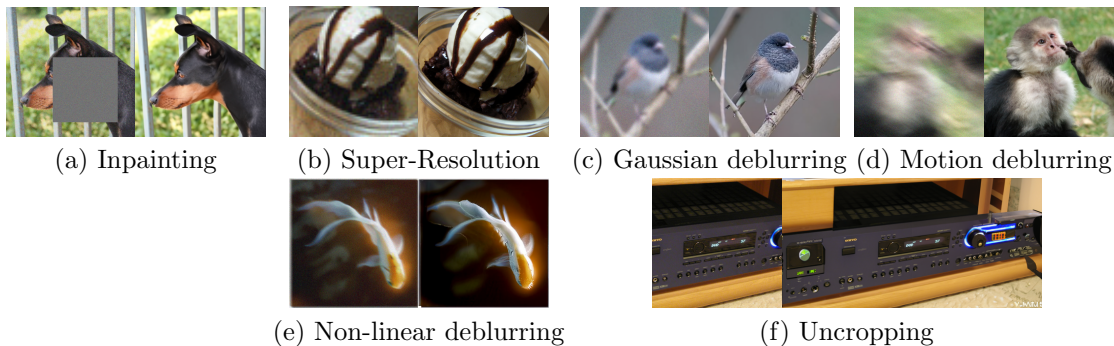


Figure 1: Examples on solving noisy image inverse problems on ImageNet validation set using our proposed method without task specific model finetuning or training.

substituting the unmasked pixels of the generated image with the noisy pixels on \mathbf{y} in every diffusion generation step. To solve a wider range of tasks such as super-resolution and deblurring, researchers further decompose \mathcal{A} using the singular value decomposition (SVD) to obtain its column and null space Song et al. (2021a); Wang et al. (2023); Kawar et al. (2022). The null space contents is refined with the help of both the pre-trained diffusion generative model and the known column space contents from \mathbf{y} . Specifically, in each diffusion generation step i , Kawar et al. (2022) fills the column space contents in \mathbf{x}_i with a noisy input image \mathbf{y} and then predict \mathbf{x}_{i-1} ; Wang et al. (2023) refines the column space of \mathbf{x}_i using both the current prediction \mathbf{x}_i and the input \mathbf{y} , and then use the refined \mathbf{x}_i to denoise and obtain \mathbf{x}_{i-1} . However, those plug-in approaches can only work for linear inverse problems, and each task requires an SVD decomposition of the operator \mathcal{A} . To solve a wider range of non-linear inverse problem when the SVD decomposition becomes impossible, another line of research directly use the conditional probability $p_i(\mathbf{y}|\mathbf{x}_i)$ to guide the generation process Chung et al. (2022, 2023); Meng and Kabashima (2022); Song et al. (2023c,a); Rout et al. (2023); Song et al. (2023b); Hu et al. (2023). Notice that the guidance score function $\nabla_{\mathbf{x}_i} \log \mathbb{E}_{p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)}[p_0(\mathbf{y}|\mathbf{x}_0)]$ is the gradient of the expected cost taken over distribution $p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$, and computing $p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$ requires run the diffusion generation process from step i . To relieve the computation burden, the diffusion posterior sampling (DPS) method Chung et al. (2023) approximates $p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$ with a Gaussian distribution $q_{0|i}(\mathbf{x}_0|\mathbf{x}_i) = \mathcal{N}(\boldsymbol{\mu}_i(\mathbf{x}_i), \sigma_i^2 \mathbf{I})$, where the mean $\boldsymbol{\mu}_i(\mathbf{x}_i)$ is the minimum mean squared error (MMSE) estimation of the clean image \mathbf{x}_0 given the current noisy image \mathbf{x}_i . The score function, denoted as $\nabla_{\mathbf{x}_i} \log \mathbb{E}_{p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)}[p_0(\mathbf{y}|\mathbf{x}_0)]$, is approximated by computing the gradient of $\log p_0(\mathbf{y}|\mathbf{x}_0)$ at sample $\mathbf{x}_0 = \boldsymbol{\mu}_i(\mathbf{x}_i)$ with the highest density mass in distribution $q_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$, i.e., $\nabla_{\mathbf{x}_i} \log \mathbb{E}_{p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)}[p_0(\mathbf{y}|\mathbf{x}_0)]$. In this

work, we further improve the estimation of DPS via the policy gradient method by using multiple samples from $q_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$.

Our Contributions. We propose a new method to estimate the guidance score for solving image inverse problems in each diffusion step, which results in a better image restoration quality and fidelity compared with the ground truth image. Our contributions are summarized as follows:

- We redefine each noisy image as a policy, where the predicted clean image serves as a state selected by the policy. DPG is a new approach for estimating the guidance score function given the input image \mathbf{y} .
- DPG eliminates the need for computing a closed-form pseudo-inverse or performing SVD decomposition. Leveraging a pre-trained diffusion generative model, our method can address a broad range of image inverse problems without requiring task-specific model fine-tuning. Selected image inverse results generated by DPG are presented in Fig. 1.
- The score function estimated by DPG is theoretically more accurate than DPS, particularly in the earlier stages of the generation process. In experiments, DPG excels in restoring high-frequency details in images. Quantitative evaluations conducted on FFHQ, ImageNet, and LSUN image restoration tasks demonstrate that our proposed method achieves improvements in both image restoration quality and consistency compared to the ground truth.

2 Methodology

2.1 Preliminaries: Diffusion Models

Let \mathbf{x}_0 be a high-fidelity image that follows an unknown distribution $p_0(\mathbf{x}_0)$. Distribution $p_0(\mathbf{x}_0)$ can be turned into a Gaussian by gradually adding Gaussian

noise by iteratively for N steps. The noise injection in each step i is as follows:

$$\mathbf{x}_i = \sqrt{1 - \beta_i} \mathbf{x}_{i-1} + \sqrt{\beta_i} \mathbf{z}_i, \mathbf{z}_i \sim \mathcal{N}(0, \mathbf{I}). \quad (2)$$

When the iterations $N \rightarrow \infty$, the evolution of $\mathbf{x}_{0:T}$ is then a result of a continuous time stochastic differential equation. Let T be a constant, denote $t_i := \frac{i}{N}T$ and $\beta(t_i) := \frac{\beta_{i+1}}{\Delta t}$, the value of each noisy image \mathbf{x}_i is the result of $\mathbf{x}(t_i)$ obtained by the following stochastic differential equation (Song et al., 2021b, Eq. (11)):

$$d\mathbf{x} = -\frac{1}{2}\beta(t)\mathbf{x}dt + \sqrt{\beta(t)}d\mathbf{w}, t \in [0, T], \quad (3)$$

where \mathbf{w} is a Wiener process. Denote $\bar{\alpha}(t) := \exp\left(-\int_0^t \beta(s)ds\right)$ and let $\sigma(t) := \sqrt{1 - \bar{\alpha}(t)}$, the conditional probability $p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0))$ is a Gaussian distribution, i.e.,

$$p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0)) = \mathcal{N}(\mathbf{x}(t)|\sqrt{\bar{\alpha}(t)}\mathbf{x}(0), (1 - \bar{\alpha}(t))\mathbf{I}). \quad (4)$$

At the final step T , $\bar{\alpha}(T) \rightarrow 0$ and therefore, $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$ is a Gaussian distribution.

Generating high resolution image \mathbf{x}_0 is equivalent to sample from distribution $p_0(\mathbf{x}_0)$. An intuitive solution is to run the reverse process of equation 3 by starting from the final step $\mathbf{x}(T) \sim \mathcal{N}(0, \mathbf{I})$. According to (Song et al., 2021b, Eq. (29)), the reverse SDE of equation 3 is as follows:

$$d\mathbf{x} = \left[-\frac{\beta(t)}{2}\mathbf{x} - \beta(t)\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) \right] dt + \sqrt{\beta(t)}d\mathbf{w}, \quad (5)$$

where $p_t(\mathbf{x}(t)) := \int p_{t|0}(\mathbf{x}(t)|\mathbf{x}_0)p(\mathbf{x}_0)d\mathbf{x}_0$ is the marginal distribution of noisy image \mathbf{x}_t .

A high quality image $\mathbf{x}_0 \sim p_0(\mathbf{x}_0)$ can be generated by numerically solving the reverse time SDE equation 5. The earliest DDPM solver interpolates $[0, T]$ with N discrete timestamps. Let $t_i := \frac{i}{N}T$, $\mathbf{x}_i := \mathbf{x}(t_i)$ be the i -th interpolated time point and value, DDPM generate \mathbf{x}_0 via the following equation:

$$\mathbf{x}_{i-1} = \frac{1}{\sqrt{\alpha_i}}\mathbf{x}_t + \frac{1 - \alpha_i}{\sqrt{1 - \alpha_i}}(\sigma_i \nabla_{\mathbf{x}} \log p_i(\mathbf{x}_i)) + \sqrt{\beta_i}\bar{\mathbf{z}}_i, \quad (6)$$

where $\bar{\mathbf{z}}_i \sim \mathcal{N}(0, \mathbf{I})$ is a random Gaussian noise sampled in step $i \in [N]$, $\beta_i := \beta(t_i)\Delta t$ and $\bar{\alpha}_i := \bar{\alpha}(t_i)$, $\sigma_i := \sqrt{1 - \bar{\alpha}_i}$. Coefficient $\alpha_{i-1} := 1 - \beta_{i-1} = \exp(-\beta_{i-1}) = \exp\left(-\int_{t_{i-1}}^{t_i} \beta(s)ds\right) = \frac{\bar{\alpha}_{i-1}}{\bar{\alpha}_i}$.

The Denoising Diffusion Implicit Model (DDIM) accelerates DDPM by generating \mathbf{x}_0 using:

$$\begin{aligned} \mathbf{x}_{i-1} &= \frac{\sqrt{\bar{\alpha}_{i-1}}}{\sqrt{\bar{\alpha}_i}}\mathbf{x}_i + \sqrt{\bar{\alpha}_{i-1}} \left(\sqrt{\frac{1 - \bar{\alpha}_{i-1}}{\bar{\alpha}_{i-1}}} - \sqrt{\frac{1 - \bar{\alpha}_i}{\bar{\alpha}_i}} \right) \\ &\quad \times (\sigma_i \log p_i(\mathbf{x}_i)). \end{aligned} \quad (7)$$

2.2 Solving Inverse Problems as Posterior Sampling

Our goal is to find the high quality image \mathbf{x}_0 that matches the input \mathbf{y} , i.e., find the root of equation 1. With the priori $p_0(\mathbf{x}_0)$, generating \mathbf{x}_0 that matches \mathbf{y} can be viewed as sampling from the following conditional probability:

$$p_0(\mathbf{x}_0|\mathbf{y}) = \frac{p_0(\mathbf{x}_0)p_0(\mathbf{y}|\mathbf{x}_0)}{p(\mathbf{y})} \propto p_0(\mathbf{x}_0)p_0(\mathbf{y}|\mathbf{x}_0). \quad (8)$$

The conditional probability $p_0(\mathbf{y}|\mathbf{x}_0)$ is a function of the reconstruction loss denoted by $\ell_{\mathbf{y}}(\mathbf{x}_0)$, i.e.,

$$p_0(\mathbf{y}|\mathbf{x}_0) \propto \exp\left(-\frac{1}{Z}\ell_{\mathbf{y}}(\mathbf{x}_0)\right), \quad (9)$$

where Z is a constant independent of \mathbf{x}_0 and \mathbf{y} . For image \mathbf{y} that is distorted by Gaussian random noise, i.e., $\mathbf{n} \sim \mathcal{N}(0, \sigma_{\mathbf{y}}^2\mathbf{I})$, the reconstruction loss $\ell_{\mathbf{y}}(\mathbf{x}_0) = \|\mathbf{y} - \mathcal{A}(\mathbf{x}_0)\|_2^2$.

We then discuss sampling from $p_0(\mathbf{x}_0|\mathbf{y})$ by running the reverse of the forward diffusion equation 3. Let $p_t(\mathbf{x}(t)|\mathbf{y}) := \int p_{t|0}(\mathbf{x}(t)|\mathbf{x}_0)p(\mathbf{x}_0|\mathbf{y})d\mathbf{x}_0$ be the marginal distribution of \mathbf{x}_t given input image \mathbf{y} . By replacing $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ in equation 5 with $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}|\mathbf{y})$, we can sample $\mathbf{x}_0 \sim p_0(\mathbf{x}_0|\mathbf{y})$ through the following SDE:

$$d\mathbf{x} = \left[-\frac{\beta(t)}{2}\mathbf{x} - \beta(t)\nabla_{\mathbf{x}} \log p_t(\mathbf{x}(t)|\mathbf{y}) \right] dt + \sqrt{\beta(t)}d\mathbf{w}. \quad (10)$$

To solve equation 10 numerically through the DDPM equation 6 or DDIM equation 7 method, we need to compute the score function of the $\mathbf{s}_i(\mathbf{x}_i, \mathbf{y}) := \nabla_{\mathbf{x}(t_i)} \log p_{t_i}(\mathbf{x}(t_i)|\mathbf{y}) = \nabla_{\mathbf{x}_i} \log p_i(\mathbf{x}_i|\mathbf{y})$, which can be decomposed by:

$$\begin{aligned} \mathbf{s}_i(\mathbf{x}_i, \mathbf{y}) &:= \sigma_i \nabla_{\mathbf{x}_i} \log p_i(\mathbf{x}_i|\mathbf{y}) = \sigma_i \nabla_{\mathbf{x}_i} \log p_i(\mathbf{x}_i, \mathbf{y}) \\ &= \underbrace{\sigma_i \nabla_{\mathbf{x}_i} \log p_i(\mathbf{x}_i)}_{\boldsymbol{\epsilon}_{\theta}(\mathbf{x}_i, i)} + \sigma_i \nabla_{\mathbf{x}_i} \log p_i(\mathbf{y}|\mathbf{x}_i). \end{aligned} \quad (11)$$

The first term $\nabla_{\mathbf{x}_i} \log p_i(\mathbf{x}_i)$ is the score function of the marginal distribution $p_i(\mathbf{x}_i)$, which is learned through the generative model parameterized by $\boldsymbol{\theta}$, i.e., $\sigma_i \nabla_{\mathbf{x}_i} \log p_i(\mathbf{x}_i) = \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_i, i)$. Computing the second term $\nabla_{\mathbf{x}_i} \log p_i(\mathbf{y}|\mathbf{x}_i)$ in equation 11 is hard because $p_i(\mathbf{y}|\mathbf{x}_i) = \int_{\mathbf{x}_0} p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)p_0(\mathbf{y}|\mathbf{x}_0)d\mathbf{x}_0$ requires integral over posterior $p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$.

2.3 Computing $\nabla_{\mathbf{x}} \log p_i(\mathbf{x}_i|\mathbf{y})$ as Policy Gradient

The computation of the gradient $\nabla_{\mathbf{x}_i} \log p_i(\mathbf{y}|\mathbf{x}_i)$ can be decomposed as follows:

$$\nabla_{\mathbf{x}_i} \log p_i(\mathbf{y}|\mathbf{x}_i)$$

$$\begin{aligned}
 &= \nabla_{\mathbf{x}_i} \left(\log \int p_{0|i}(\mathbf{x}_0|\mathbf{x}_i) p_0(\mathbf{y}|\mathbf{x}_0) d\mathbf{x}_0 \right) \\
 &= \frac{1}{\int p_{0|i}(\mathbf{x}_0|\mathbf{x}_i) p_0(\mathbf{y}|\mathbf{x}_0) d\mathbf{x}_0} \\
 &\quad \times \left(\nabla_{\mathbf{x}_i} \int p_{0|i}(\mathbf{x}_0|\mathbf{x}_i) p_0(\mathbf{y}|\mathbf{x}_0) d\mathbf{x}_0 \right) \\
 &\propto \nabla_{\mathbf{x}_i} \int \underbrace{p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)}_{\text{State Density Function}} \underbrace{p_0(\mathbf{y}|\mathbf{x}_0)}_{\text{Cost}} d\mathbf{x}_0 =: \tilde{\mathbf{s}}_i(\mathbf{x}_i, \mathbf{y}).
 \end{aligned} \tag{12}$$

Notice that $\tilde{\mathbf{s}}_i(\mathbf{x}_i, \mathbf{y})$ contains the directional information about the guidance score $\nabla_{\mathbf{x}_i} \log p_i(\mathbf{y}|\mathbf{x}_i)$. To compute $\tilde{\mathbf{s}}_i(\mathbf{x}_i, \mathbf{y})$, notice that the generated image \mathbf{x}_0 is determined by the intermediate noisy image \mathbf{x}_i , and the conditional probability $p_0(\mathbf{y}|\mathbf{x}_0)$ can be viewed as a cost of \mathbf{x}_0 . Therefore, the computation of the score function $\tilde{\mathbf{s}}_i(\mathbf{x}_i, \mathbf{y})$ in equation 12 is closely related to policy gradient in reinforcement learning, where $p_{0|t}(\mathbf{x}_0|\mathbf{x}_t)$ is the state density function by choosing policy \mathbf{x}_i , and $p_0(\mathbf{y}|\mathbf{x}_0)$ is the cost. The following theorem enables us to compute the score function equation 12 from the policy gradient perspective:

Theorem 1 (Leibniz Rule) *Suppose $p_0(\mathbf{x}_0)$ is the probability measure of $N_{train}, N_{train} < \infty$ high quality training images. Then for all $i \in [N]$, we can compute the score function $\tilde{\mathbf{s}}_i(\mathbf{x}_i, \mathbf{y})$ from equation 12 as follows:*

$$\tilde{\mathbf{s}}_i(\mathbf{x}_i, \mathbf{y}) = \mathbb{E}_{p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)} [p_0(\mathbf{y}|\mathbf{x}_0) \nabla_{\mathbf{x}_i} \log p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)]. \tag{13}$$

Notice that the diffusion model is usually trained with a finite number of images, Theorem 1 works for a wide range of pretrained diffusion models. Proof of Theorem 1 is provided in Appendix 5.1. A benefit for using equation 13 is that we do not need to compute the derivative of $p_0(\mathbf{y}|\mathbf{x}_0)$ (i.e., the reconstruction loss $\ell_{\mathbf{y}}(\mathbf{x}_0)$) and therefore do not require $p_0(\mathbf{y}|\mathbf{x}_0)$ to be differentiable.

2.4 Implementation Details

2.4.1 Tractable Monte Carlo sampling

Computing the score function in equation 13 requires sampling clean images \mathbf{x}_0 from a complex distribution $p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$ and then compute its density. To facilitate the sampling and computation, similar to Chung et al. (2023); Song et al. (2023c), we select a Gaussian distribution $q_{0|i}(\mathbf{x}_0|\mathbf{x}_i) = \mathcal{N}(\boldsymbol{\mu}_i(\mathbf{x}_i), r_i^2 \mathbf{I})$ to approximate $p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$. The mean $\boldsymbol{\mu}_i(\mathbf{x}_i)$ and variance r_i for $q_{0|i}$ should be selected so that $q_{0|i}$ is close to $p_{0|i}$, i.e., the KL divergence between $p_{0|i}$ and $q_{0|i}$ is small. The following Lemma provides us with a good mean and variance selection principle:

Lemma 1 *For any distribution $p_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$ whose density function is absolutely continuous on \mathbb{R}^{d_x} , the optimum Gaussian distribution $q_{0|i}(\mathbf{x}_0|\mathbf{x}_i) = \mathcal{N}(\boldsymbol{\mu}_i(\mathbf{x}_i), r_i^2 \mathbf{I})$ that minimizes the KL divergence $D(p_{0|i}(\mathbf{x}_0|\mathbf{x}_i) || q_{0|i}(\mathbf{x}_0|\mathbf{x}_i))$ is:*

$$\boldsymbol{\mu}_i(\mathbf{x}_i) = \mathbb{E}_{p_{0|i}}[\mathbf{x}_0|\mathbf{x}_i], r_i^2 = \frac{1}{d_x} \mathbb{E}_{p_{0|i}} [\|\mathbf{x}_0 - \boldsymbol{\mu}_i(\mathbf{x}_i)\|_2^2]. \tag{14}$$

We then select the mean $\boldsymbol{\mu}_i(\mathbf{x}_i)$ and variance r_i using Lemma 1.

(1). $\boldsymbol{\mu}_i(\mathbf{x}_i)$: According to equation 14, the optimum mean $\boldsymbol{\mu}_i(\mathbf{x}_i)$ is the MMSE estimation of \mathbf{x}_0 given \mathbf{x}_i , which requires an integral over $p_{0|i}(\mathbf{x}_0|\mathbf{x}_i) = \frac{p_0(\mathbf{x}_0) p_{i|0}(\mathbf{x}_i|\mathbf{x}_0)}{p_i(\mathbf{x}_i)}$. The Tweedie’s estimator Efron (2011); Kim and Ye (2021) is capable of estimating $\boldsymbol{\mu}_i(\mathbf{x}_i)$ using the conditional $p_{i|0}(\mathbf{x}_i|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_i; \sqrt{\bar{\alpha}_i} \mathbf{x}_0, \sqrt{1 - \bar{\alpha}_i} \mathbf{I})$ and the score function $\nabla_{\mathbf{x}_i} \log p_i(\mathbf{x}_i) = \frac{\boldsymbol{\epsilon}_{\theta}(\mathbf{x}_i, i)}{\sqrt{1 - \bar{\alpha}_i}}$, i.e.,

$$\begin{aligned}
 &\boldsymbol{\mu}_i(\mathbf{x}_i) \\
 &:= \sqrt{\frac{1 - \bar{\alpha}_i}{\bar{\alpha}_i}} \left(\frac{1}{\sqrt{1 - \bar{\alpha}_i}} \mathbf{x}_i - \sqrt{1 - \bar{\alpha}_i} \nabla_{\mathbf{x}_i} \log p_i(\mathbf{x}_i) \right) \\
 &\stackrel{(a)}{=} \frac{1}{\sqrt{\bar{\alpha}_i}} (\mathbf{x}_i - \sqrt{1 - \bar{\alpha}_i} \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_i, i)),
 \end{aligned} \tag{15}$$

where equality (a) is obtained because $\boldsymbol{\epsilon}_{\theta}(\mathbf{x}_i, i) = \sqrt{1 - \bar{\alpha}_i} \nabla_{\mathbf{x}_i} \log p_i(\mathbf{x}_i)$.

(2). r_i : According to equation 14, r_i should be the average pixel-wise prediction error of the clean image \mathbf{x}_0 . However, computing r_i using the Tweedie’s formula requires computing the Hessian of $\log p_i(\mathbf{x}_i)$, which is computational challenging in diffusion models. As an alternative, notice that the reconstruction error between the input \mathbf{y} and the current prediction $\mathcal{A}(\boldsymbol{\mu}_i(\mathbf{x}_i))$, i.e., $\|\mathbf{y} - \mathcal{A}(\boldsymbol{\mu}_i(\mathbf{x}_i))\|_2$ can also reflect the prediction error of the clean image. A small reconstruction error indicates our prediction $\boldsymbol{\mu}_i(\mathbf{x}_i)$ is accurate and the error $\|\mathbf{x}_0 - \boldsymbol{\mu}_i(\mathbf{x}_i)\|_2^2$ should be small as well. Therefore, we intuitively select

$$r_i = \sqrt{\frac{1}{C \times H \times W} \|\mathbf{y} - \mathcal{A}(\boldsymbol{\mu}_i(\mathbf{x}_i))\|_2^2}, \tag{16}$$

where C, H, W are the channels, height and width of the image \mathbf{x}_0 .

With the distribution $q_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$, we can estimate the score function $\tilde{\mathbf{s}}_i(\mathbf{x}_i, i)$ inequation 12 using the Monte Carlo method. Let $\mathbf{x}_0^{(m)} \sim q_{0|i}(\mathbf{x}_0|\mathbf{x}_i), m = 1, \dots, N_{mc}$ samples drawn i.i.d from distribution $q_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$. Following equation 13, the $\tilde{\mathbf{s}}_i(\mathbf{x}_i, i)$ can be computed by:

$$\begin{aligned}
 &\tilde{\mathbf{s}}_i(\mathbf{x}_i, \mathbf{y}) \\
 &\approx \mathbb{E}_{q_{0|i}(\mathbf{x}_0|\mathbf{x}_i)} [p_0(\mathbf{y}|\mathbf{x}_0) \nabla_{\mathbf{x}_i} \log q_{0|i}(\mathbf{x}_0|\mathbf{x}_i)]
 \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E}_{q_{0|i}(\mathbf{x}_0|\mathbf{x}_i)} \left[p_0(\mathbf{y}|\mathbf{x}_0) \nabla_{\mathbf{x}_i} \left(-\frac{1}{2r_i^2} \|\mathbf{x}_0 - \boldsymbol{\mu}_i(\mathbf{x}_i)\|_2^2 \right) \right] \\
 &= -\frac{1}{2r_i^2 N_{\text{mc}}} \sum_{m=1}^{N_{\text{mc}}} \left(p_0(\mathbf{y}|\mathbf{x}_0^{(m)}) \nabla_{\mathbf{x}_i} \|\mathbf{x}_0 - \boldsymbol{\mu}_i(\mathbf{x}_i)\|_2^2 \right)
 \end{aligned} \tag{17}$$

2.4.2 Selection of $p_0(\mathbf{y}|\mathbf{x}_0)$

After sample $\mathbf{x}_0^{(m)}, m \in [N_{\text{mc}}]$ from distribution $q_{0|i}(\mathbf{x}_0|\mathbf{x}_i) = \mathcal{N}(\boldsymbol{\mu}_i(\mathbf{x}_i), r_i^2 \mathbf{I})$, we need to compute the conditional probability $p_0(\mathbf{y}|\mathbf{x}_0^{(m)})$ for each data sample. Ideally, the conditional distribution $p_0(\mathbf{y}|\mathbf{x}_0^{(m)}) \propto \exp\left(-\frac{1}{2}\ell_{\mathbf{y}}(\mathbf{x}_0^{(m)})\right)$ should be a fixed distribution for all the generation steps i . However, in the initial steps of the diffusion generation process (i is large), the prediction $\boldsymbol{\mu}_i(\mathbf{x}_i)$ is far away from the target image, the reconstruction loss $\ell_{\mathbf{y}}(\mathbf{x}_0)$ between the input image \mathbf{y} and the posterior sampled image $\mathcal{A}(\mathbf{x}_0^{(m)})$ can be very large. If Z is too small, the conditional probability $p_0(\mathbf{y}|\mathbf{x}_0)$ is always 0 for all the samples $\mathbf{x}_0^{(m)}$ in step i ; while in later stage of the diffusion generation process (i is small), the reconstruction loss $\ell_{\mathbf{y}}(\mathcal{A}(\boldsymbol{\mu}_i(\mathbf{x}_i)))$ can be small, selecting a large Z may lead to the fact that all the conditional probabilities of $\mathbf{x}_0^{(m)} \sim q_{0|i}(\mathbf{x}_0|\mathbf{x}_i)$ is close to one. Selecting an invariant Z for all the diffusion generation steps lead to the that all the conditional probability $p_0(\mathbf{y}|\mathbf{x}_0^{(m)})$ are too close to 0 or 1, and the score function $\tilde{s}_i(\mathbf{x}_i, \mathbf{y})$ is not accurate. To prevent these phenomena, we choose Z_i adaptively and assign a different normalizing factor for the conditional probability $p_0(\mathbf{y}|\mathbf{x}_0^{(m)})$ in each generation step i . For Gaussian noise, we select $p_0(\mathbf{y}|\mathbf{x}_0) \propto \exp\left(-\frac{1}{Z_i} \|\mathbf{y} - \mathcal{A}(\mathbf{x}_0)\|_2^2\right)$, where each $Z_i = r_i^2 = \frac{1}{C \times H \times W} \|\mathbf{y} - \mathcal{A}(\boldsymbol{\mu}_i(\mathbf{x}_i))\|_2^2$, i.e., $p(\mathbf{y}|\mathbf{x}_0)$ has the same variance as $q(\mathbf{x}_0|\mathbf{x}_i)$ defined in equation 16; for Poisson noise, we select the conditional distribution $p_0(\mathbf{y}|\mathbf{x}_0) \propto \exp\left(-\frac{1}{Z_i} \|\mathbf{y} - \mathcal{A}(\mathbf{x}_0)\|_1\right)$ to be an exponential distribution, where $Z_i = \|\mathbf{y} - \mathcal{A}(\boldsymbol{\mu}_i(\mathbf{x}_i))\|_1$.

2.4.3 Reward Shaping

Similar to policy gradient in reinforcement learning, direct MC estimation of the policy gradient equation 13 suffers from high estimation variance. To reduce the estimation variance, we leverage the reward shaping technique Ng et al. (1999). The goal is to deduct the a bias term $b_i(\mathbf{x}_i) = \mathbb{E}_{q_{0|i}(\mathbf{x}_0|\mathbf{x}_i)}[p_0(\mathbf{y}|\mathbf{x}_0)]$ from the cost $p_0(\mathbf{y}|\mathbf{x}_0)$ of each sample \mathbf{x}_0 . Notice that $b_i(\mathbf{x}_i)$ should be a scalar independent of samples \mathbf{x}_0 , therefore we deduct a different bias term $b^{(m)}$ from each sample $p_0(\mathbf{y}|\mathbf{x}_0^{(m)})$ using the leave-one-out cross-

validation, i.e.,

$$b_i^{(m)} := \frac{1}{N_{\text{mc}} - 1} \sum_{j=1, j \neq m}^{N_{\text{mc}}} p_0(\mathbf{y}|\mathbf{x}_0^{(j)}). \tag{18}$$

We can then improve the MC estimation from equation 17 by:

$$\begin{aligned}
 \tilde{s}_i(\mathbf{x}_i, \mathbf{y}) &= \\
 &= -\frac{\sum_{m=1}^{N_{\text{mc}}} (p_0(\mathbf{y}|\mathbf{x}_0^{(m)}) - b_i^{(m)}) \nabla_{\mathbf{x}_i} \|\mathbf{x}_0^{(m)} - \boldsymbol{\mu}_i(\mathbf{x}_i)\|_2^2}{2r_i^2 N_{\text{mc}}}.
 \end{aligned} \tag{19}$$

2.4.4 Score Function Re-scaling

Notice that the score function computed from equation 19 contains only direction information. The exact norm of the gradient $\nabla_{\mathbf{x}_i} \log p_i(\mathbf{y}|\mathbf{x}_i)$ is unknown. We need to re-scale the computed score function $\tilde{s}_i(\mathbf{x}_i, \mathbf{y})$ with norm B , i.e., assume that $\sigma_i \nabla_{\mathbf{x}_i} \log p_i(\mathbf{y}|\mathbf{x}_i) \approx B \cdot \frac{1}{\|\tilde{s}_i(\mathbf{x}_i, \mathbf{y})\|_2} \tilde{s}_i(\mathbf{x}_i, \mathbf{y})$ and plug it into equation 11 to compute the score function $\mathbf{s}_i(\mathbf{x}_i, \mathbf{y})$, i.e.,

$$\mathbf{s}_i(\mathbf{x}_i, \mathbf{y}) \approx \epsilon_{\theta}(\mathbf{x}_i, i) + B \cdot \frac{\tilde{s}_i(\mathbf{x}_i, \mathbf{y})}{\|\tilde{s}_i(\mathbf{x}_i, \mathbf{y})\|_2}. \tag{20}$$

2.5 Algorithm Description

Using the score function estimated by equation 20, we can solve the image inversion problems with the conditional guided diffusion using the standard DDPM/DDIM sampling method. The image restoration restore the target image \mathbf{x}_0 is displayed in Algorithm 1.

2.6 Connection with Diffusion Posterior Sampling (DPS)

In this part, we discuss the relationship between the score function obtained by DPS Chung et al. (2023) and our proposed method in Corollary 1.

Corollary 1

When $r_i \rightarrow 0$ and assume we have infinite number of MC samples, if $p_0(\mathbf{y}|\mathbf{x}_0)$ is a Gaussian distribution, then the score function 17 is approximately

$$\tilde{s}_i(\mathbf{x}_i, \mathbf{y}) = \frac{1}{2\sigma_{\mathbf{y}}^2 r_i} p_0(\mathbf{y}|\mathbf{x}_0) \nabla_{\mathbf{x}_i} \ell_{\mathbf{y}}(\mathbf{x}_i), \tag{21}$$

whose direction is the same of the score function in DPS Chung et al. (2023).

Proof of Corollary 1 is provided in Appendix 5.2. Corollary 1 shows that the score function $\tilde{s}_i(\mathbf{x}_i, \mathbf{y}) \approx$




Figure 2: Evolution of the reconstruction loss $\|\mathbf{y} - \mathcal{A}(\boldsymbol{\mu}_i(\mathbf{x}_i))\|_2$ of the DPS and DPG method.

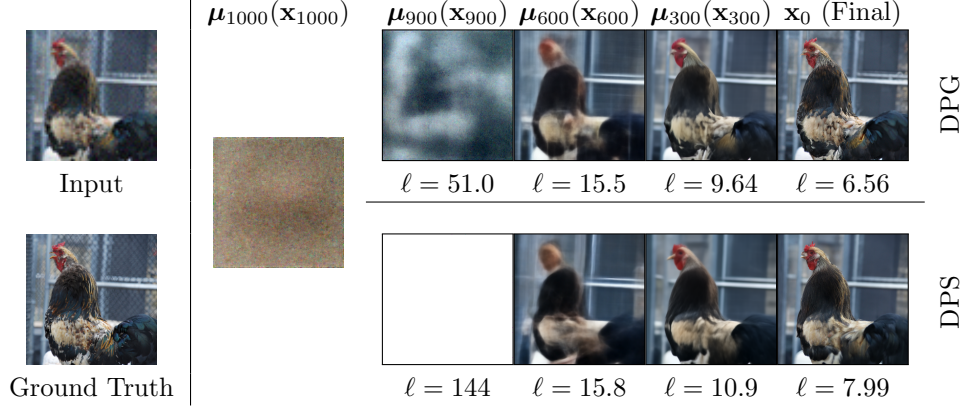


Figure 3: Image generation procedure and the reconstruction loss $\ell := \|\mathbf{y} - \mathcal{A}(\boldsymbol{\mu}_i(\mathbf{x}_i))\|_2$ by using DPG and DPS methods in super-resolution.

Algorithm 1 Diffusion Policy Gradient (DPG)

Require: Number of generation steps N , input image \mathbf{y} , reconstruction loss function $\ell_{\mathbf{y}}(\mathcal{A}(\cdot))$

$\mathbf{x}_N \sim \mathcal{N}(0, \mathbf{I})$

for $i = N$ to 1 **do**

$\boldsymbol{\mu}_i(\mathbf{x}_i) \leftarrow \frac{1}{\sqrt{\alpha_i}} (\mathbf{x}_i + \sqrt{1 - \alpha_i} \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_i, i))$ ▷
 equation 14.

$r_i \leftarrow \sqrt{\frac{1}{C \times H \times W} \|\mathbf{y} - \mathcal{A}(\boldsymbol{\mu}_i(\mathbf{x}_i))\|_2^2}$ ▷ equation 16

$\mathbf{x}_0^{(m)} \leftarrow \mathcal{N}(\boldsymbol{\mu}_i(\mathbf{x}_i), r_i^2 \mathbf{I}), m = 1, \dots, N_{\text{mc}},$ ▷

Sample from $q(\mathbf{x}_0 | \mathbf{x}_i)$

$b_i^{(m)} \leftarrow \frac{1}{N_{\text{mc}} - 1} \sum_{j=1, j \neq m}^{N_{\text{mc}}} p_0(\mathbf{y} | \mathbf{x}_0^{(j)})$ ▷ Reward
 Shaping equation 18

$\tilde{\mathbf{s}}_i(\mathbf{x}_i, \mathbf{y}) \leftarrow - \frac{\sum_{m=1}^{N_{\text{mc}}} (p_0(\mathbf{y} | \mathbf{x}_0^{(m)}) - b_i^{(m)}) \nabla_{\mathbf{x}_i} \|\mathbf{x}_0^{(m)} - \boldsymbol{\mu}_i(\mathbf{x}_i)\|_2^2}{2r_i^2 N_{\text{mc}}}$

▷ Diffusion Policy Gradient equation 19

$\mathbf{s}_i(\mathbf{x}_i, \mathbf{y}) \leftarrow \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_i, i) + B \frac{\tilde{\mathbf{s}}_i(\mathbf{x}_i, \mathbf{y})}{\|\tilde{\mathbf{s}}_i(\mathbf{x}_i, \mathbf{y})\|_2}$

$\mathbf{x}_{i-1} \leftarrow \frac{1}{\sqrt{\alpha_i}} \mathbf{x}_i + \frac{1 - \alpha_i}{\sqrt{1 - \alpha_i}} \mathbf{s}_i(\mathbf{x}_i, \mathbf{y}) + \sqrt{\beta_i} \bar{\mathbf{z}}_i.$ ▷

DDPM

Or $\mathbf{x}_{i-1} \leftarrow \sqrt{\frac{\alpha_{i-1}}{\alpha_i}} \mathbf{x}_i + \sqrt{\frac{1 - \alpha_{i-1}}{\alpha_i}} \left(\sqrt{\frac{1 - \alpha_{i-1}}{\alpha_i}} - \sqrt{\frac{1 - \alpha_i}{\alpha_i}} \right) \mathbf{s}_i(\mathbf{x}_i, \mathbf{y})$ ▷ DDIM

end for

Return image \mathbf{x}_0

$\nabla_{\mathbf{x}_i} \ell_{\mathbf{y}}(\mathbf{x}_i)$ by DPS Chung et al. (2023) is accurate when $r_i \rightarrow 0$, i.e., in later-stages of the diffusion generation process. However, in initial stages of the image generation (i.e., i is large), the score function obtained by DPS is inaccurate. As is shown in Fig. 2, the reconstruction loss $\ell_{\mathbf{y}}(\mathbf{x}_i)$ by running the DPS algorithm in the initial image generation stages (i.e., $i \geq 750$) larger compared with our proposed DPG method. Fig. 3 plots the intermediate recovered figures during the diffusion process. Since DPG has a more accurate estimation of the guidance score function, the shape and the sketch of the image is recovered at an earlier stage compared with the DPS method (i.e., at step $i = 900$, noisy image generated by DPG has the sketch of the

chicken, while the image generated by DPS is blank.) More results for deblurring experiments can be found in the Appendix.

3 Experiments

In this section, we show experimental results obtained by our DPG method and compare its performance with other state-of-the-art methods.

3.1 Quantitative Results on Noisy Linear Inverse Problems

Experiment Setup Similar to Chung et al. (2023); Song et al. (2023b), we test the performance of our proposed algorithm on three datasets: the FFHQ 256×256 dataset Karras et al. (2019), the ImageNet dataset Deng et al. (2009) and the LSUN-Bedroom dataset Yu et al. (2015). We consider four types of image inverse tasks: (1) Inpainting, where a size of 128×128 mask is added randomly on the input image; (2) 4×super-resolution with average pooling; (3) Gaussian deblurring with kernel size 61×61 and standard deviation of 3.0; (4) Motion deblurring with kernel size of 61 and intensity value 0.5 generated by¹. We consider that the input image is noisy, i.e., Gaussian noise with variance $\sigma_{\mathbf{y}} = 0.05$ or Poisson noise with rate $\lambda = 1.0$ is added on the input image. For FFHQ experiments, we use the pretrained model from Chung et al. (2023) (trained on 4.9k images on FFHQ) and test the performance of 1k validation set; For Imagenet experiments, we use the unconditional Imagenet 256×256 generation model from Dhariwal and Nichol (2021) and the 1k images are selected from ²; For LSUN experiments, we use the pretrained LSUN-Bedroom model from Dhariwal and Nichol (2021) and test the performance on LSUN-Bedroom validation set containing 300 images. All experiments are run on an NVIDIA A100 GPU.

¹<https://github.com/LeviBorodenko/motionblur>

²https://github.com/XingangPan/deep-generative-prior/blob/master/scripts/imagenet_val_1k.txt




Figure 4: Results on solving linear noisy inverse problems (inpainting, super-resolution and Gaussian deblurring) on ImageNet Dataset. The input image is distorted by random Gaussian noise $\sigma_y = 0.05$.




Figure 5: Image Restoration Results for Motion Deblurring on ImageNet 256×256 .




Figure 6: Image Restoration Results on ImageNet with Poisson Noise $\lambda = 1.0$.

We compare the performance with the following methods: Denoising Diffusion Null Space models (DDNM+) Wang et al. (2023) for noisy problems, Diffusion Posterior Sampling (DPS) Chung et al. (2022), Denoising Diffusion Restoration Models (DDRM) Kawar et al. (2022) and the Reddif Mardani et al. (2023). The key parameters for different methods are displayed in Appendix 6. For LSUN dataset, we also report the results obtained by the DDIM sampler with 200 steps generation.

Evaluation Metrics We measure both the image restoration quality and consistency compared with the ground-truth image. For image restoration quality, we compute the Fréchet inception distance between the restored images and the ground truth images; For image restoration consistency, we compute the LPIPS score Zhang et al. (2018) (VGG Net) and the Peak Signal to Noise Ratio (PSNR) between the restored image and the ground truth image. Quantitative evaluation results are displayed in Table 3.1. Selected image restoration samples when the observation noise are Gaussian and Poisson are displayed in Fig. 4 and Fig. 6.

Analysis The FID and LPIPS score of our proposed DPG method is smaller than DPS method, indicating that DPG has a better image restoration quality than DPS method. This is because the estimation of the score function by DPG is more accurate than DPS,

Table 1: Quantitative Results on Inpainting, Super-Resolution and Gaussian deblurring on FFHQ, Imagenet and LSUN-Bedroom (256 × 256) Datasets

Method	Inpainting (Random)			Super-Resolution (4×)			deblurring (Gauss)			deblurring (Motion)		
	FID↓	LPIPS↓	PSNR↑	FID↓	LPIPS↓	PSNR↑	FID↓	LPIPS↓	PSNR↑	FID↓	LPIPS↓	PSNR↑
FFHQ 1k Validation Set												
DPG (DDPM)	22.44	0.181	22.17	22.49	0.214	26.61	22.29	0.216	26.02	24.44	0.223	26.38
DPS (DDPM) Chung et al. (2023)	33.12	0.216	21.83	39.35	0.214	25.67	<u>44.05</u>	<u>0.257</u>	24.93	39.02	0.242	24.92
DDRM Kawar et al. (2022)	27.47	0.172	23.44	62.15	0.294	25.36	74.92	0.332	23.36	N/A	N/A	N/A
DDNM+ Wang et al. (2023)	27.34	0.173	23.29	46.13	0.260	27.41	63.19	0.301	27.70	N/A	N/A	N/A
ImageNet 1k Validation Set												
DPG (DDPM)	41.86	0.258	17.41	31.02	0.293	22.91	34.43	0.314	22.10	36.15	0.343	21.67
DPS (DDPM) Chung et al. (2023)	45.95	0.267	17.69	43.60	0.340	23.10	54.76	0.386	20.04	56.08	0.386	20.55
DDRM Kawar et al. (2022)	50.94	0.246	19.13	51.77	0.355	24.17	72.49	0.345	22.62	N/A	N/A	N/A
DDNM+ Wang et al. (2023)	50.50	0.246	19.16	51.08	0.362	24.00	71.74	0.410	24.90	N/A	N/A	N/A
RED-Diff Mardani et al. (2023)	192.96	0.292	20.03	74.39	0.434	23.39	62.79	0.380	23.65	N/A	N/A	N/A
LSUN-Bedroom Validation Set												
DPG (DDPM)	34.32	0.218	18.90	31.44	0.262	23.48	38.72	0.277	22.39	34.44	0.284	22.82
DPG (DDIM)	34.39	0.209	19.74	33.86	0.269	23.99	46.08	0.322	22.25	45.08	0.355	21.97
DPS Chung et al. (2023)	35.91	0.218	18.58	37.42	0.284	23.67	48.10	0.320	22.25	50.09	0.358	21.73
DDRM Kawar et al. (2022)	37.61	0.205	19.59	50.96	0.310	24.10	59.04	0.353	22.64	N/A	N/A	N/A
DDNM+ Wang et al. (2023)	37.03	0.204	19.55	50.15	0.296	24.48	74.40	0.336	24.38	N/A	N/A	N/A
RED-Diff Mardani et al. (2023)	44.35	0.240	20.49	75.76	0.380	24.67	64.70	<u>0.314</u>	25.27	N/A	N/A	N/A

Table 2: Quantitative Results on Non-linear deblurring on Imagenet (256 × 256) Datasets

Method	FID↓	LPIPS↓	PSNR↑
DPG (Ours)	88.15	0.464	20.73
DPS	120.79	0.484	17.79

especially in the initial stages of the diffusion generation process. Therefore, the shape and structure of the image can be recovered in an earlier stage of the diffusion process, this gives room to recover high frequency details in later stage of the image generation. Notice that DDNM+ and DDRM uses a plug-in estimation, i.e., the known pixels in \mathbf{y} are directly used in the generation process. Therefore, the PSNR of DDNM+ and DDRM are better than the proposed DPG method, but the recovered high frequency detailed features are less than our proposed DPG method. DPG has a smaller LPIPS score and FID score in most tasks. The Rediff method is not robust to input noise, the quality of the restored image degrades significantly. Moreover, since DPG has a better estimation of the guidance score, we can combine DPG with the DDIM solver. Results on LSUN dataset shows there is little quality loss compared with the 1000 steps DDPM generation.

3.2 Quantitative Experiments on Non-Linear Image Inverse Problems

Our method does not require the operator \mathcal{A} to be linear, and thus can be applied to non-linear image inversion problems. We consider the input image is distorted by the non-linear blur kernel Tran et al. (2021), and compare the image restoration results of our DPG method with the DPS Chung et al. (2023). Selected image restoration results are displayed in Fig. 7 and quantitative evaluations are displayed in Table 3.2. According to Fig. 7, DPG can restore more details of the original image, and this results in a smaller LPIPS score.



Figure 7: Image Restoration Results for Non-linear Deblurring on ImageNet.

4 Conclusions and Limitations

In this paper, we proposed a new method to estimate the guidance score function to solve image inverse problems using pre-trained diffusion generative models. Our method is robust when the input image is perturbed by random noise, and can be used for solving non-linear inverse problems. Experiments demonstrate that the proposed method can improve image restoration quality in both human eye evaluation and quantitative metrics for a wide range of tasks such as inpainting, super-resolution and deblurring.

While the score function computed by DPG using equation 20 eliminates the necessity for a differentiable loss function $\ell_{\mathbf{y}}(\mathbf{x}_0)$, we observe, in the specific case of JPEG restoration, that the quality of the generated image is less stable and far worse compared to the results achieved by IIGDM Song et al. (2023b).

Acknowledging the superiority of our method over alternatives like DPS and DDNM, it is essential to note that the generation speed of our approach DPG is slower when the operator becomes complex (i.e., the non-linear deblurring kernel). Moreover, current deblurring experiments are carried out when the blur kernel is known. In the future, we would like to explore solving deblurring problems with unknown kernels and apply the method on latent diffusion models Rombach et al. (2021).

References

- Avrahami, O., Fried, O., and Lischinski, D. (2023). Blended latent diffusion. *ACM Trans. Graph.*, 42(4).
- Avrahami, O., Lischinski, D., and Fried, O. (2022). Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18208–18218.
- Blattmann, A., Rombach, R., Oktay, K., and Ommer, B. (2022). Retrieval-augmented diffusion models.
- Chung, H., Kim, J., Mccann, M. T., Klasky, M. L., and Ye, J. C. (2023). Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*.
- Chung, H., Sim, B., Ryu, D., and Ye, J. C. (2022). Improving diffusion models for inverse problems using manifold constraints. In *Advances in Neural Information Processing Systems*.
- Couairon, G., Verbeek, J., Schwenk, H., and Cord, M. (2023). Diffedit: Diffusion-based semantic image editing with mask guidance. In *The Eleventh International Conference on Learning Representations*.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- Dhariwal, P. and Nichol, A. (2021). Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794.
- Efron, B. (2011). Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614.
- Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851.
- Ho, J. and Salimans, T. (2021). Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*.
- Hu, Y., Wang, Y., and Zhang, J. (2023). Deargan: Degradation-aware face restoration with gan prior. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(9):4603–4615.
- 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. In *Advances in Neural Information Processing Systems*.
- Kim, K. and Ye, J. C. (2021). Noise2score: Tweedie’s approach to self-supervised image denoising without clean images. In *Advances in Neural Information Processing Systems*.
- 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*.
- Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.-Y., and Ermon, S. (2022). SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*.
- Meng, X. and Kabashima, Y. (2022). Diffusion model based posterior sampling for noisy linear inverse problems. *arXiv preprint arXiv:2211.12343*.
- Ng, A. Y., Harada, D., and Russell, S. J. (1999). Policy invariance under reward transformations: Theory and application to reward shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning, ICML ’99*, page 278–287, San Francisco, CA, USA.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2021). High-resolution image synthesis with latent diffusion models.
- Rout, L., Raoof, N., Daras, G., Caramanis, C., Dimakis, A. G., and Shakkottai, S. (2023). Solving linear inverse problems provably via posterior sampling with latent diffusion models. *arXiv preprint arXiv:2307.00619*.
- 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. PMLR.
- 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.
- Song, J., Vahdat, A., Mardani, M., and Kautz, J. (2023b). 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. (2023c). Loss-guided diffusion models for plug-and-play controllable generation. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 32483–32498. PMLR.
- Song, Y., Shen, L., Xing, L., and Ermon, S. (2021a). Solving inverse problems in medical imaging with score-based generative models. In *NeurIPS 2021 Workshop on Deep Learning and Inverse Problems*.

- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. (2021b). Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*.
- Tran, P., Tran, A. T., Phung, Q., and Hoai, M. (2021). Explore image deblurring via encoded blur kernel space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11956–11965.
- Wang, Y., Yu, J., and Zhang, J. (2023). Zero-shot image restoration using denoising diffusion null-space model. *The Eleventh International Conference on Learning Representations*.
- Weber, A. G. (2006). The usc-sipi image database: Version 5. <http://sipi.usc.edu/database/>.
- Yu, F., Zhang, Y., Song, S., Seff, A., and Xiao, J. (2015). Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. (2018). The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*.