

# CONSISTENCY-GUIDED REVERSE SAMPLING FOR GENERAL LINEAR INVERSE PROBLEMS

Anonymous authors  
 Paper under double-blind review

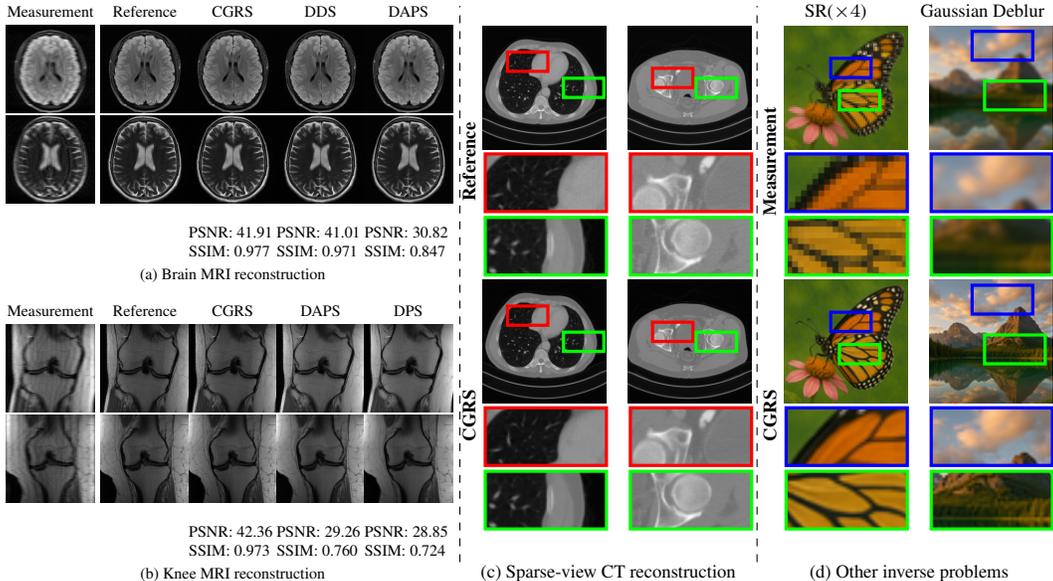


Figure 1: Overview of Consistency-Guided Reverse Sampling (CGRS). Our method introduces a flexible framework for solving a wide range of inverse problems by incorporating consistency constraints into the reverse sampling process. In (a)(b), we show visual comparison on brain and knee MRI reconstruction tasks under highly undersampled conditions, highlighting the ability of our CGRS to preserve anatomical structures and fine details. In (c), we present the performance of our CGRS for sparse-view CT reconstruction, demonstrating improved edge sharpness and reduced artifacts. In (d), we illustrate the performance of our CGRS for classical natural image restoration tasks including super-resolution and Gaussian deblurring on ImageNet, showing strong generalization across different inverse problems.

## ABSTRACT

Diffusion models have recently demonstrated strong potential in solving inverse problems by leveraging the priors of learned data. However, the standard reverse sampling process often struggles to correct errors in the early stages, leading to suboptimal reconstructions, particularly in ill-posed or underdetermined settings. To overcome this issue, we present Consistency-Guided Reverse Sampling (CGRS), a novel framework that integrates measurement-consistency optimization into each reverse diffusion step. CGRS refines the intermediate denoised estimates by solving a linear least-squares problem, thereby improving consistency with the observed measurements. This correction mechanism mitigates error accumulation along the sampling trajectory and enhances overall reconstruction fidelity. Furthermore, CGRS naturally supports flexible acceleration by allowing the reverse process to start from a coarse optimization-based reconstruction, effectively reducing the number of reverse steps with negligible degradation in reconstruction quality. Experimental results in various linear inverse problems

demonstrate that CGRS consistently yields superior performance. Our code is available at <https://anonymous.4open.science/r/diffusion-inverse-sampling-00D1>.

## 1 INTRODUCTION

Inverse problems exist in many scientific and engineering fields, including computational imaging (Boys et al., 2023; Afonso et al., 2010; Sun & Simon, 2021), medical imaging Suetens (2017); Ravishankar et al. (2019); Sun et al. (2023), and remote sensing (Liu et al., 2021; 2022b). These problems aim to recover an unknown signal  $\mathbf{x}_0$  from noisy and incomplete measurements  $\mathbf{y} = \mathcal{A}(\mathbf{x}_0) + \mathbf{n}$ . However, due to the inherent ill-posedness and noise in the forward process, the solution is often ambiguous and unstable.

Recent advances in diffusion models (DMs) (Ho et al., 2020; Song et al., 2020b; Dhariwal & Nichol, 2021; Rombach et al., 2022) have enabled strong generative capabilities by learning complex data distributions via a noise-to-data transformation defined by stochastic differential equations (SDEs). When used as priors in Bayesian inverse problems (Meng & Kabashima, 2022; Chung et al., 2022a; Boys et al., 2023), DMs allow sampling from the posterior distribution by modifying the reverse SDE to include measurement information.

Several works have adopted this idea to solve inverse problems through conditional sampling (Chung et al., 2022a; Song et al., 2023; Chung et al., 2022b), replacing the unconditional score with a conditional one. Despite their effectiveness, existing methods often require hundreds of denoising steps and still suffer from accumulated error and weak measurement consistency.

To address these challenges, we propose *Consistency-Guided Reverse Sampling (CGRS)*, a novel framework that integrates measurement-consistency optimization into each reverse diffusion step. Instead of relying solely on score-based updates, CGRS explicitly enforces consistency with the measurement by projecting intermediate samples back onto the feasible space defined by the forward model. Moreover, we introduce an efficient sampling variant, Fast CGRS, which leverages a coarse optimization-based reconstruction to initialize the reverse process from a lower noise level. This significantly reduces the number of required diffusion steps while maintaining high-quality reconstruction.

We validate CGRS across a range of linear inverse problems, including MRI reconstruction, sparse-view CT reconstruction, and classical image restoration tasks. Experimental results demonstrate that CGRS consistently improves reconstruction fidelity and efficiency compared to state-of-the-art methods, highlighting its robustness and generalizability.

## 2 PRELIMINARIES

### 2.1 DIFFUSION MODELS

Assume that we have a  $d$ -dimensional random variable  $\mathbf{x}_0 \in \mathbb{R}^d$  with an unknown distribution  $q_0(\mathbf{x}_0)$ . Diffusion models (DMs) (Ho et al., 2020; Song et al., 2020b; Kingma et al., 2021; Dhariwal & Nichol, 2021; Rombach et al., 2022) usually define a forward process  $\{\mathbf{x}_t\}_{t \in [0, T]}$  with  $T > 0$  starting with  $\mathbf{x}_0$ , such that for any  $t \in [0, T]$ , the distribution of  $\mathbf{x}_t$  conditioned on  $\mathbf{x}_0$  satisfies

$$q_{0t}(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t | \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}). \quad (1)$$

The forward process can usually be represented as  $\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}$  ( $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$ ), where  $\alpha_t \in \mathbb{R}^+$  is a differentiable function of  $t$  with bounded derivative,  $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$ , and the choice for  $\alpha_t$  is referred to as the noise schedule of a DM. (Kingma et al., 2021) have proved that the following stochastic differential equation (SDE) has the same transition distribution  $q_{0t}(\mathbf{x}_t | \mathbf{x}_0)$  as in Equation 1 for any  $t \in [0, T]$ :

$$d\mathbf{x}_t = f(\mathbf{x}_t, t)dt + g(t)d\mathbf{w}_t, \quad \mathbf{x}_0 \sim q_0(\mathbf{x}_0), \quad (2)$$

where  $\mathbf{w}_t \in \mathbb{R}^d$  is the standard Wiener process,  $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$  and  $g : \mathbb{R} \rightarrow \mathbb{R}$  are the drift and diffusion coefficient that can take different functional forms (e.g., Variance-Preserving SDEs (VP-SDEs) and Variance-Exploding SDEs (VE-SDEs) in (Song et al., 2020b)).

Under some regularity conditions, (Song et al., 2020b) show that the forward process in Equation 2 has an equivalent reverse process from time  $T$  to 0, starting with the marginal distribution  $q_T(\mathbf{x}_T)$ :

$$\begin{aligned} d\mathbf{x}_t &= [f(\mathbf{x}_t, t) - g^2(t)\nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t)]dt + g(t)d\mathbf{w}_t, \\ \mathbf{x}_T &\sim q_T(\mathbf{x}_T), \end{aligned} \quad (3)$$

where  $\mathbf{w}_t \in \mathbb{R}^d$  is the standard Wiener process in the reverse process. The only unknown term in Equation 3 is the score function  $\nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t)$  at each time  $t$ . In practice, DMs use a neural network  $\epsilon_\theta(\mathbf{x}_t, t)$  ( $\mathbf{s}_\theta(\mathbf{x}_t, t)$ ) parameterized by  $\theta$  to estimate the scaled score function:  $-\sqrt{1 - \bar{\alpha}_t}\nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t)$  ( $\nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t)$ ). The parameter  $\theta$  is optimized by minimizing the following objective (Ho et al., 2020; Song et al., 2020b):

$$\min_{\theta} \mathbb{E}_{q_0(\mathbf{x}_0), \epsilon, t} [\omega(t) \|\epsilon_\theta(\mathbf{x}_t, t) - \epsilon\|_2^2], \quad (4)$$

or

$$\min_{\theta} \mathbb{E}_{q_t(\mathbf{x}_t), t} [\omega(t) \|\mathbf{s}_\theta(\mathbf{x}_t, t) + \sqrt{1 - \bar{\alpha}_t}\nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t)\|_2^2], \quad (5)$$

where  $\mathbf{x}_t \sim q_t(\mathbf{x}_t)$ ,  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ,  $t \sim \mathcal{U}([0, 1])$ , and  $\omega(t) > 0$  is a weighting function. Once the parameterized scaled score function  $\epsilon_\theta$  is obtained, we can approximate the reverse diffusion process in Equation 3 and simulate it using a sampling algorithm. A common choice is the DDPM sampling scheme (Ho et al., 2020), which updates the sample as follows:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left[ \mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{x}_t, t) \right] + \sqrt{1 - \bar{\alpha}_t} \mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (6)$$

where  $\alpha_t$  is the noise schedule and  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ .

In practice, the reverse process is not limited to the DDPM update rule, and other numerical solvers (e.g., DDIM (Song et al., 2020a), higher-order ODE/SDE solvers (Lu et al., 2022; Liu et al., 2022a; Zheng et al., 2023)) can also be adopted. More generally, the discrete-time reverse diffusion step can be expressed as

$$\mathbf{x}_{t-1} = h(\mathbf{x}_t, t; \epsilon_\theta) + \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{w}}), \quad (7)$$

where  $h : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is a deterministic update function parameterized by the score function, and  $\mathbf{w}_t$  is the injected Gaussian noise with covariance  $\Sigma_{\mathbf{w}}$ . Equation 7 thus characterizes the general non-linear dynamics of a discrete-time diffusion sampler.

This general formulation also makes it straightforward to incorporate additional constraints—such as enforcing measurement consistency—into the sampling trajectory, which forms the basis of our proposed Consistency-Guided Reverse Sampling (CGRS) method.

## 2.2 DIFFUSION MODELS FOR SOLVING INVERSE PROBLEMS

Inverse problems, arising from a wide range of applications in various domains, including computational imaging, medical imaging, and remote sensing, among others, aim to recover unknown signals from partial measurements with potential noise. The forward process is usually lossy, resulting in an ill-posed signal recovery task where a unique solution does not exist. In general, a forward model takes the form of

$$\mathbf{y} = \mathcal{A}(\mathbf{x}_0) + \mathbf{n}, \quad (8)$$

where  $\mathcal{A}(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^m$  is the measurement function,  $\mathbf{x}_0$  represents the original data,  $\mathbf{y} \in \mathbb{R}^m$  is the observed measurement, and  $\mathbf{n}$  symbolizes the noise in the measurement process, often modeled as  $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \beta_{\mathbf{y}}^2 \mathbf{I})$ .

In the Bayesian framework, given measurements  $\mathbf{y} \in \mathbb{R}^m$  from the forward model Equation 8, one utilizes  $q(\mathbf{x})$  as the prior, and samples from the posterior distribution  $q(\mathbf{x} | \mathbf{y})$ , where the relationship is formally established based on the Bayes' rule:  $q(\mathbf{x} | \mathbf{y}) \propto q(\mathbf{y} | \mathbf{x})q(\mathbf{x})$ . Leveraging the diffusion model as the prior, we can solve inverse problems by replacing the score function in Equation 3 with the conditional score function  $\nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | \mathbf{y})$ :

$$d\mathbf{x}_t = [f(\mathbf{x}_t, t) - g(t)^2 \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | \mathbf{y})] dt + g(t)d\mathbf{w}_t, \quad (9)$$

where, according to Bayes' rule, we can get:

$$\nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | \mathbf{y}) = \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log q_t(\mathbf{y} | \mathbf{x}_t). \quad (10)$$

Then, by integrating Equation 10, Equation 9 can generally be written as

$$d\mathbf{x}_t = [f(\mathbf{x}_t, t) - g(t)^2[\nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log q_t(\mathbf{y} | \mathbf{x}_t)]]dt + g(t)d\mathbf{w}_t. \quad (11)$$

Many papers, including classifier-guidance diffusion (Song et al., 2020b; Dhariwal & Nichol, 2021), DPS (Chung et al., 2022a), IIGDM (Song et al., 2023), MCG (Chung et al., 2022b), MPGD (He et al., 2023), FreeDoM (Yu et al., 2023), DSG (Yang et al., 2024), attempt to leverage pre-trained diffusion models to solve Equation 11 for various inverse problems.

## 3 METHOD

### 3.1 CONSISTENCY-GUIDED REVERSE SAMPLING

---

#### **Algorithm 1** Consistency-Guided Reverse Sampling (CGRS)

---

**Require:** Model  $\epsilon_\theta$ , Measurement  $\mathbf{y}$ , Measurement function  $\mathcal{A}$ , Noise schedule  $\{\alpha_t\}_{t=1}^T$ .

- 1: Sample  $\mathbf{x}_T \sim \mathcal{N}(0, (1 - \bar{\alpha}_T)\mathbf{I})$  ▷Initial noise vector
- 2: **for**  $t = T, T - 1, \dots, 1$  **do**
- 3:    $\hat{\epsilon}_t = \epsilon_\theta(\mathbf{x}_t, t)$  ▷Compute the score
- 4:    $\hat{\mathbf{x}}_0(\mathbf{x}_t) = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t}\hat{\epsilon}_t)$  ▷Predict  $\hat{\mathbf{x}}_0$  using Tweedie’s formula
- 5:    $\hat{\mathbf{x}}_0(\mathbf{x}_t, \mathbf{y}) \in \arg \min_{\hat{\mathbf{x}}_0} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0(\mathbf{x}_t))\|^2$ , ▷Solve with initial point  $\hat{\mathbf{x}}_0(\mathbf{x}_t)$
- 6:   Sample  $\mathbf{x}_{t-1} \sim \mathcal{N}(\sqrt{\bar{\alpha}_{t-1}}\hat{\mathbf{x}}_0(\mathbf{x}_t, \mathbf{y}), (1 - \bar{\alpha}_{t-1})\mathbf{I})$  ▷Sample  $\mathbf{x}_{t-1}$
- 7: **end for**
- 8: **return**  $\mathbf{x}_0$

---

When applying diffusion models (DMs) to inverse problems, Equation 7 can be reformulated as

$$\mathbf{x}_{t-1} = h(\mathbf{x}_t, \mathbf{y}, t; \epsilon_\theta) + \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \Sigma_w), \quad (12)$$

which corresponds to drawing samples from the conditional distribution  $q_t(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{y})$ .

For a reverse step  $t \rightarrow t - 1$  using the Euler–Maruyama discretization, a common and accurate strategy is to introduce  $\mathbf{x}_0$  as a latent variable and approximate the update through a plug-in posterior. This procedure starts from the identity

$$q_t(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{y}) = \int q_t(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)q_t(\mathbf{x}_0 | \mathbf{x}_t, \mathbf{y}) d\mathbf{x}_0. \quad (13)$$

In practice, direct evaluation of this integral is generally intractable. We first examine the posterior distribution  $q_t(\mathbf{x}_0 | \mathbf{x}_t, \mathbf{y})$ . By Bayes’ rule, it can be expressed as

$$q_t(\mathbf{x}_0 | \mathbf{x}_t, \mathbf{y}) \propto p(\mathbf{y} | \mathbf{x}_0) q_t(\mathbf{x}_0 | \mathbf{x}_t). \quad (14)$$

Assuming a Gaussian likelihood, we obtain

$$p(\mathbf{y} | \mathbf{x}_0) \propto \exp\left(-\frac{1}{2\beta^2}\|\mathbf{y} - \mathcal{A}(\mathbf{x}_0)\|_2^2\right) \quad (15)$$

Prior works on diffusion posterior sampling (Chung et al., 2022a; Song et al., 2023; Chung et al., 2022b) approximate  $q_t(\mathbf{x}_0 | \mathbf{x}_t)$  by a Gaussian of the form

$$q_t(\mathbf{x}_0 | \mathbf{x}_t) \approx \mathcal{N}(\mathbf{x}_0; \hat{\mathbf{x}}_0(\mathbf{x}_t), \frac{1 - \bar{\alpha}_t}{\bar{\alpha}_t}\mathbf{I}). \quad (16)$$

Building on Equation 16, the posterior mean can be expressed in closed form, as stated in Proposition 1, by applying Tweedie’s formula (Efron, 2011; Kim & Ye, 2021).

**Proposition 1** For the case of VP-SDE or VE-SDE sampling (Song et al., 2020b),  $q_t(\mathbf{x}_0 | \mathbf{x}_t)$  has the unique posterior mean at

$$\hat{\mathbf{x}}_0(\mathbf{x}_t) = \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \frac{1}{\sqrt{\bar{\alpha}_t}}[\mathbf{x}_t + (1 - \bar{\alpha}_t)\mathbf{s}_\theta(\mathbf{x}_t, t)]. \quad (17)$$

For the case of DDPM sampling (Ho et al., 2020), the formula can be rewritten as

$$\hat{\mathbf{x}}_0(\mathbf{x}_t) = \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \frac{1}{\sqrt{\bar{\alpha}_t}}[\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)]. \quad (18)$$

The posterior mean can be computed as the solution of

$$\hat{\mathbf{x}}_0(\mathbf{x}_t, \mathbf{y}) = \arg \min_{\mathbf{x}_0} \frac{1}{2\beta_{\mathbf{y}}^2} \|\mathbf{y} - \mathcal{A}(\mathbf{x}_0(\mathbf{x}_t))\|_2^2 + \frac{\bar{\alpha}_t}{2(1 - \bar{\alpha}_t)} \|\mathbf{x}_0 - \hat{\mathbf{x}}_0(\mathbf{x}_t)\|_2^2 \quad (19)$$

At late times (small noise), the influence of the prior diminishes and the update reduces to the simplified correction

$$\hat{\mathbf{x}}_0(\mathbf{x}_t, \mathbf{y}) \approx \arg \min_{\hat{\mathbf{x}}_0} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0(\mathbf{x}_t))\|^2, \quad \text{initialized at } \hat{\mathbf{x}}_0(\mathbf{x}_t), \quad (20)$$

which is equivalent to a measurement-corrected refinement of the denoised estimate. We can use many classical approaches to solve this optimization problem, including the Conjugate Gradient (CG) (Nazareth, 2008), the Alternating Direction Method of Multipliers (ADMM) (Gabay & Mercier, 1976), the Trust Region method (Conn et al., 2000).

In unconditional sampling, the posterior distribution  $q_t(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$  is often employed since the posterior mean combines information from both the noisy sample  $\mathbf{x}_t$  and the network-based estimate of  $\mathbf{x}_0$ , which is usually represented as

$$q_t(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}\left(\tilde{\mu}(\mathbf{x}_t, \mathbf{x}_0), \tilde{\beta}_t \mathbf{I}\right), \quad (21)$$

where  $\tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0) = \frac{\sqrt{\bar{\alpha}_{t-1}}(1 - \alpha_t)}{1 - \bar{\alpha}_t} \mathbf{x}_0 + \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \mathbf{x}_t, \tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} (1 - \alpha_t).$

This formulation reduces variance and stabilizes the reverse diffusion trajectory, which is essential when the estimate of  $\mathbf{x}_0$  is purely inferred from the generative prior.

However, in conditional sampling scenarios, the situation differs significantly. Once the estimate  $\hat{\mathbf{x}}_0(\mathbf{x}_t, \mathbf{y})$  has been refined through measurement guidance, it already incorporates the external information  $\mathbf{y}$ . In this case, relying on the bridge distribution can introduce inconsistency, as its mean  $\tilde{\mu}(\mathbf{x}_t, \mathbf{x}_0)$  couples the condition-informed estimate  $\hat{\mathbf{x}}_t$  with the unconditional latent  $\mathbf{x}_t$ . This mixture may reintroduce artifacts from the unconditional path and dilute the influence of the measurement constraints.

**Proposition 2** In conditional diffusion sampling, once the estimate  $\hat{\mathbf{x}}_0$  incorporates measurement information  $\mathbf{y}$ , the posterior distribution  $q_t(\mathbf{x}_{t-1} | \mathbf{x}_t, \hat{\mathbf{x}}_0)$  can be well approximated by the marginal distribution  $q_t(\mathbf{x}_{t-1} | \hat{\mathbf{x}}_0)$ .

As shown in Proposition 2, we use the marginal distribution

$$q_t(\mathbf{x}_{t-1} | \hat{\mathbf{x}}_0) = \mathcal{N}(\sqrt{\bar{\alpha}_{t-1}}\hat{\mathbf{x}}_0, (1 - \bar{\alpha}_{t-1})\mathbf{I}), \quad (22)$$

which treats the condition-informed  $\hat{\mathbf{x}}_0$  as a surrogate ground-truth signal. This design avoids dependence on the potentially inconsistent  $\hat{\mathbf{x}}_t$ , ensuring that each reverse step propagates information aligned with the observation  $\mathbf{y}$ . As a result, the ancestral form often yields more stable and measurement-consistent reconstructions in conditional diffusion sampling. More details can be found in Appendix A.

Overall, the reverse diffusion process can be expanded into a three-step procedure: (1) sampling  $\hat{\mathbf{x}}_0(\mathbf{x}_t) \sim q_t(\mathbf{x}_0 | \mathbf{x}_t)$ , (2) solving  $\hat{\mathbf{x}}_0(\mathbf{x}_t, \mathbf{y}) \in \arg \min_{\hat{\mathbf{x}}_0} \frac{1}{2} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|_2^2$  with initial point  $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ , and (3) sampling  $\mathbf{x}_{t-1} \sim \mathcal{N}(\sqrt{\bar{\alpha}_{t-1}}\hat{\mathbf{x}}_0(\mathbf{x}_t, \mathbf{y}), (1 - \bar{\alpha}_{t-1})\mathbf{I})$ . We repeat this process, gradually reducing noise until  $\mathbf{x}_0$  is sampled. We refer to this process as *consistency-guided reverse sampling*. The full procedure is summarized in Algorithm 1.

**Algorithm 2** Fast CGRS

---

**Require:** Model  $\epsilon_\theta$ , Measurement  $\mathbf{y}$ , Measurement function  $\mathcal{A}$ , Noise schedule  $\{\alpha_t\}_{t=1}^S$ .

- 1: Initialize vector  $\hat{\mathbf{x}}_0 = \mathbf{0}$
- 2:  $\hat{\mathbf{x}}_0 \in \arg \min_{\hat{\mathbf{x}}_0} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|^2$  ▷Initialize  $\hat{\mathbf{x}}_0$  with a coarse reconstruction
- 3:  $\mathbf{x}_S \sim \mathcal{N}(\sqrt{\bar{\alpha}_S}\hat{\mathbf{x}}_0, (1 - \bar{\alpha}_S)\mathbf{I})$  ▷Sample  $\mathbf{x}_S$
- 4: **for**  $t = S, S - 1, \dots, 1$  **do**
- 5:    $\hat{\epsilon}_t = \epsilon_\theta(\mathbf{x}_t, t)$  ▷Compute the score
- 6:    $\hat{\mathbf{x}}_0(\mathbf{x}_t) = \frac{1}{\sqrt{\bar{\alpha}_t}} (\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t}\hat{\epsilon}_t)$  ▷Predict  $\hat{\mathbf{x}}_0$  using Tweedie’s formula
- 7:    $\mathbf{x}_0(\mathbf{x}_t) \in \arg \min_{\hat{\mathbf{x}}_0} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0(\mathbf{x}_t))\|^2$ , ▷Solve with initial point  $\hat{\mathbf{x}}_0(\mathbf{x}_t)$
- 8:   Sample  $\mathbf{x}_{t-1} \sim \mathcal{N}(\sqrt{\bar{\alpha}_{t-1}}\mathbf{x}_0(\mathbf{x}_t), (1 - \bar{\alpha}_{t-1})\mathbf{I})$  ▷Sample  $\mathbf{x}_{t-1}$
- 9: **end for**
- 10: **return**  $\mathbf{x}_0$

---

## 3.2 CGRS FOR FAST SAMPLING

Although diffusion models are known for their superior performance and generalisation capacity, their slow inference time is a critical drawback. Over the years, significant efforts have been made to alleviate this issue, including the development of improved ODE/SDE solvers (Kim et al., 2021; Liu et al., 2022a; Lu et al., 2022; Zheng et al., 2023; Lu et al., 2025) and the adoption of recursive distillation approaches (Salimans & Ho, 2022; Meng et al., 2023). However, most studies have focused on unconditional sampling, leaving conditional sampling acceleration rather under-explored.

A method called *Come-Closer-Diffuse-Faster (CCDF)* (Chung et al., 2022c) first perturbs the initial estimate by applying the forward diffusion process up to a selected time step  $t_S < T$ , where  $t_S$  marks the starting point of the reverse diffusion. Notably, this forward diffusion incurs almost no computational cost, as it does not require any neural network evaluations. We found that the *CCDF* ideas can be effectively incorporated into the *CGRS*.

Unlike *CCDF*, which samples  $\mathbf{x}_S$  at time  $t_S$  directly from the measurement  $\mathbf{y}$  via  $\mathbf{x}_S \sim \mathcal{N}(\sqrt{\bar{\alpha}_S}\mathbf{y}, (1 - \bar{\alpha}_S)\mathbf{I})$ , our approach replaces  $\mathbf{y}$  with a coarse reconstruction obtained by solving the optimization problem  $\arg \min_{\hat{\mathbf{x}}_0} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|^2$ . Specifically, we perform the following steps:

$$\begin{aligned}
 &\text{Initialize vector } \hat{\mathbf{x}}_0 = \mathbf{0}, \\
 &\hat{\mathbf{x}}_0 \in \arg \min_{\hat{\mathbf{x}}_0} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|^2, \\
 &\mathbf{x}_S \sim \mathcal{N}(\sqrt{\bar{\alpha}_S}\hat{\mathbf{x}}_0, (1 - \bar{\alpha}_S)\mathbf{I}).
 \end{aligned} \tag{23}$$

This process usually requires hundreds of iterations, which requires several seconds to complete. The full procedure is summarized in Algorithm 2.

## 4 EXPERIMENTS

## 4.1 EXPERIMENTAL SETUP

We evaluate our method using the pre-trained diffusion models provided by (Chung et al., 2023). Specifically, three pixel-space models are used, each trained on a different dataset: the fastMRI brain dataset (Zbontar et al., 2018), the fastMRI knee dataset (Zbontar et al., 2018), and the AAPM LDCT dataset (Moen et al., 2021).

**Datasets and metrics.** Building on the previously introduced models, we evaluate our method on the fastMRI and AAPM LDCT datasets. Specifically, we use 1,000 images from the fastMRI brain test set, 1,000 images from the fastMRI knee validation set, and 500 images from the AAPM LDCT validation set. Our primary evaluation metrics include peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al., 2018). For both our method and the baselines, we use the implementations provided by TorchMetrics (Nicki Skafta Detlefsen et al., 2022), with all images normalized to the [0, 1] range.

**Problem setting.** We have the following general measurement model

$$\mathbf{y} = \mathcal{M}\mathcal{T}\mathcal{S}(\mathbf{x}) = \mathcal{A}(\mathbf{x}), \mathbf{x} \in \mathbb{R}^d, \mathbf{y} \in \mathbb{R}^m, \mathcal{A} \in \mathbb{R}^{d \times m}, \quad (24)$$

where  $\mathcal{M}$  is the sub-sampling matrix,  $\mathcal{T}$  is the discrete transform matrix (i.e. Fourier, Radon), and  $\mathcal{S} = \mathbf{I}$  when we have a single-array measurement including CT, and  $\mathcal{S} = [\mathcal{S}^{(1)}, \dots, \mathcal{S}^{(c)}]$  when we have a  $c$ -coil MR parallel imaging measurement.

**Baselines.** For MRI reconstruction, we compare our CGRS method with several diffusion models. These include DPS (Chung et al., 2022a), DAPS (Zhang et al., 2025), and DDS (Chung et al., 2023), which are designed for general image restoration tasks, as well as ScoreMRI (Chung & Ye, 2022) and CSGM (Jalal et al., 2021), which are specifically tailored for MRI reconstruction using diffusion models. All experiments are conducted using equispaced sub-sampling masks. For CT reconstruction, we benchmark our method against DPS, MCG (Chung et al., 2022b), and DDS. In addition to these general-purpose inverse problem solvers, we also compare with two representative baselines: (1) Filtered Backprojection (FBP), a conventional analytic reconstruction technique, and (2) FBP-UNet (Jin et al., 2017), a supervised deep learning model that refines FBP outputs using a UNet architecture.

## 4.2 MAIN RESULTS

Task	Method	$\times 4$ Acceleration			$\times 8$ Acceleration			$\times 12$ Acceleration		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Brain-MRI	CSGM(NFE = 2311)	33.73	0.788	0.072	30.52	0.776	0.113	27.64	0.757	0.146
	ScoreMRI(NFE = 1000)	28.46	0.719	0.128	25.79	0.693	0.142	24.33	0.681	0.157
	DPS(NFE = 1000)	29.34	0.739	0.117	26.97	0.711	0.133	26.50	0.701	0.136
	DAPS(NFE = 1000)	31.69	0.829	0.058	28.30	0.778	0.094	27.11	0.753	0.108
	DDS(NFE = 1000)	40.49	0.957	0.009	34.06	0.893	0.033	31.25	0.853	0.055
	CGRS(NFE = 1000)	<b>42.51</b>	<b>0.972</b>	<b>0.006</b>	<b>34.44</b>	<b>0.906</b>	<b>0.029</b>	<b>31.37</b>	<b>0.859</b>	<b>0.051</b>
Knee-MRI	CSGM(NFE = 2311)	33.24	0.774	0.118	30.26	0.752	0.148	27.18	0.741	0.183
	ScoreMRI(NFE = 1000)	27.75	0.689	0.156	25.13	0.664	0.173	23.89	0.642	0.198
	DPS(NFE = 1000)	28.92	0.729	0.146	26.81	0.684	0.166	26.26	0.666	0.169
	DAPS(NFE = 1000)	30.94	0.779	0.106	28.07	0.709	0.152	26.93	0.676	0.165
	DDS(NFE = 1000)	39.67	0.945	0.017	33.99	0.868	0.053	31.13	0.814	0.085
	CGRS(NFE = 1000)	<b>41.73</b>	<b>0.965</b>	<b>0.011</b>	<b>34.31</b>	<b>0.886</b>	<b>0.042</b>	<b>31.25</b>	<b>0.836</b>	<b>0.078</b>

Table 1: Quantitative results of MRI reconstruction on fastMRI data. Performance comparison on brain and knee MRI reconstruction with different methods across different acceleration factors ( $\times 4$ ,  $\times 8$ ,  $\times 12$ ). NFE: number of function evaluations. The best and second-best results within each type of task are indicated in bold and underlined, respectively.

We present the quantitative results for brain and knee MRI reconstruction tasks under various acceleration settings in Table 1. As shown, our method, CGRS, consistently outperforms existing baselines across all configurations, achieving superior reconstruction accuracy and perceptual quality. Notably, CGRS maintains superior and stable performance for both brain and knee MRI reconstructions, even under high acceleration factors. Our model not only preserves global anatomical structures but also faithfully restores subtle textures and fine-grained details that are often lost by competing methods.

While other approaches tend to suffer from notable performance degradation as the acceleration factor increases—manifesting in blurred structures and perceptual artifacts—CGRS demonstrates exceptional stability. It excels at preserving high-frequency information and ensuring visual consistency across varying levels of measurement sparsity. This consistent performance across different anatomical regions and sampling conditions highlights the versatility and robustness of CGRS in addressing a broad range of MRI inverse problems.

Table 2 summarizes the quantitative results for CT reconstruction under sparse-view settings. As shown, CGRS consistently achieves the best performance across all the evaluation metrics. Compared to traditional baselines such as FBP and its UNet-enhanced variant, CGRS delivers reconstructions with significantly improved image fidelity and structural integrity. Compared to more advanced diffusion models including DPS and MCG, CGRS demonstrates apparent advantages since these methods often struggle to recover fine-grained details or suffer from diminished perceptual quality under sparse measurements. In contrast, CGRS strikes an effective balance between reconstruction accuracy and visual realism, achieving leading results across both pixel-level and perceptual evaluations. Additional visual results are provided in Appendix B to further highlight the effectiveness of CGRS.

378  
379  
380  
381  
382  
383  
384  
385

Method	18 Views		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
FBP	14.14	0.346	0.522
FBP-UNet	30.24	0.878	0.042
DPS	28.56	0.820	0.059
MCG	27.47	0.778	0.116
DDS	32.04	0.912	0.018
CGRS	<b>34.68</b>	<b>0.928</b>	<b>0.012</b>

386  
387  
388  
389

Table 2: Quantitative evaluation on AAPM data. Performance comparison on CT reconstruction with different methods across a parallel-beam geometry using 18 projection angles equally distributed across 180 degrees using the torch-radon package (Ronchetti, 2020). Bold indicates the best results.

390

#### 4.3 FURTHER RESULTS ON NATURAL IMAGE

391

392

393

394

395

396

397

398

399

400

401

402

403

404

405

406

407

408

409

410

411

412

413

414

415

416

417

418

419

420

421

422

423

424

425

426

427

428

429

430

431

Table 3: Quantitative evaluation of super-resolution ( $\times 4$ ) on FFHQ data.

We further evaluate the effectiveness of CGRS on the FFHQ (Karras et al., 2019) and ImageNet (Deng et al., 2009) datasets for classical inverse problems, including super-resolution and Gaussian deblurring. For the Gaussian deblurring task, we use a kernel of size  $61 \times 61$  with a standard deviation of 3. In the super-resolution task, we employ a bicubic resizer. We leverage pre-trained diffusion models: one trained on the FFHQ dataset by (Chung et al., 2022a), and another pre-trained model from (Dhariwal & Nichol, 2021) on the ImageNet dataset. We compare our methods with the following methods: DAPS, DDRM (Kawar et al., 2022), DPS, DDNM (Wang et al., 2022), Diff-OC (Li & Pereira, 2024) for super-resolution and Gaussian deblurring experiments. Table 3 shows that CGRS outperforms the existing diffusion-based methods significantly on FFHQ dataset. We include more qualitative results in Appendix B.

#### 4.4 ABLATION STUDY

As illustrated in Figure 2, Fast CGRS achieves significantly higher PSNR than the existing baselines under all sampling budgets. This is primarily attributed to the coarse initialization obtained via solving an optimization problem  $\arg \min_{\hat{x}_0} \|\mathbf{y} - \mathcal{A}(\hat{x}_0)\|^2$  first, which effectively narrows the sampling space and accelerates convergence. Unlike conventional methods that sample from Gaussian noise, our Fast CGRS method benefits from a structured prior of MR image reconstructed by a fast coarse reconstruction, significantly narrowing the sampling space and accelerating convergence.

422

423

424

425

426

427

428

429

430

431

	Super-Resolution ( $\times 4$ )		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Fast CGRS (NFE/S = 10)	26.59	0.783	0.375
Fast CGRS (NFE/S = 50)	30.59	0.875	0.344
Fast CGRS (NFE/S = 100)	31.25	0.886	0.342
Fast CGRS (NFE/S = 200)	31.28	0.887	0.338
Fast CGRS (NFE/S = 300)	31.32	0.889	0.337
CGRS (NFE = 1000)	32.15	0.896	0.332

Table 4: Quantitative evaluation of different initialization steps on 50 FFHQ images.

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

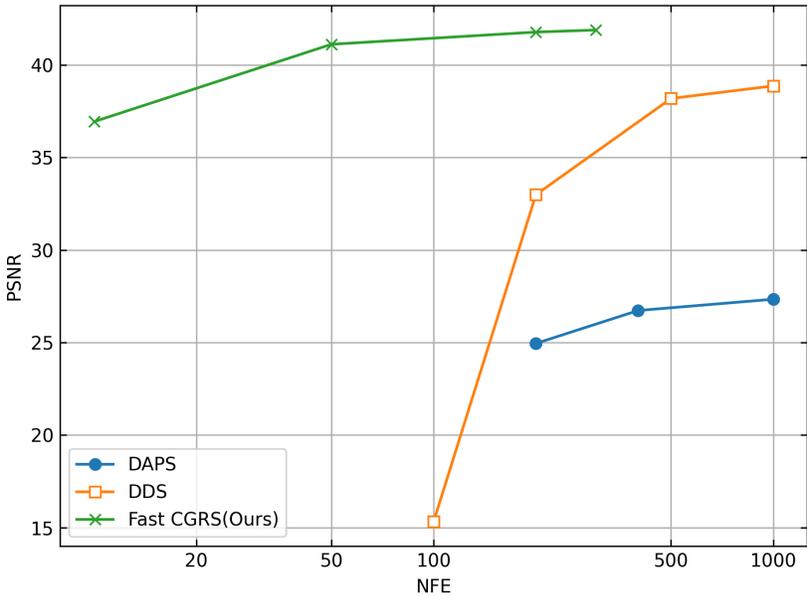


Figure 2: MRI reconstruction evaluated by PSNR vs. NFE (log scale) with equispaced sub-sampling  $\times 4$ .

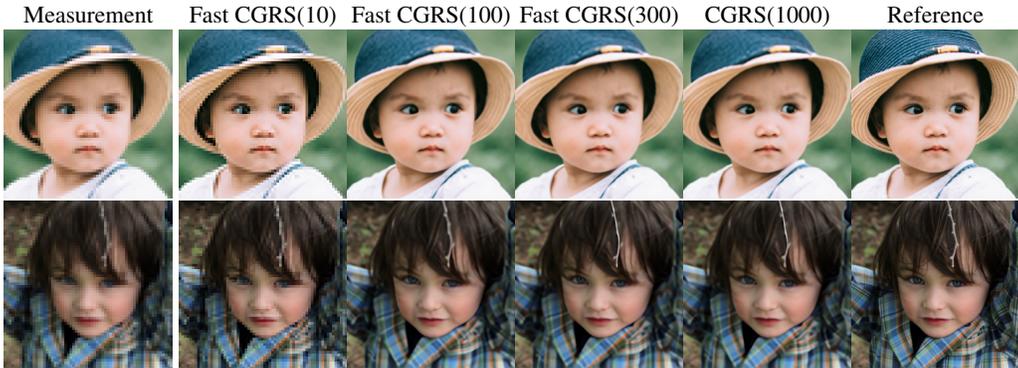


Figure 3: Visualization of different initialization steps on FFHQ dataset.

To further investigate the effect of the initialization procedure, we study the impact of selecting different starting points  $S$  for the reverse diffusion process. Table 4 reports the quantitative performance when varying  $S$ . The results show that larger  $S$  values (i.e., starting the reverse process from a later diffusion step) generally improve PSNR and SSIM while reducing LPIPS. However, the improvements plateau beyond a certain point, indicating diminishing returns. This trade-off between speed and reconstruction quality suggests that an intermediate range of NFE (e.g., 100–200) provides a balanced choice in practice, offering both computational efficiency and reliable visual fidelity. The visualized results are shown in Figure 3.

## 5 CONCLUSION

In summary, we propose Consistency-Guided Reverse Sampling (CGRS) for solving general linear inverse problems. Our method enforces measurement consistency at each step of the reverse diffusion sampling trajectory by integrating optimization-based correction, thereby correcting early-stage errors and improving data fidelity. This design also supports flexible acceleration by initializing the reverse process from a coarse optimization-based reconstruction, reducing the number of sampling steps without compromising quality. In the experiments, we demonstrate that CGRS generates samples with improved reconstruction accuracy, stability, and efficiency compared to the existing diffusion-based methods across a wide range of challenging linear inverse problems.

## REFERENCES

- 486  
487  
488 Manyá V Afonso, José M Bioucas-Dias, and Mário AT Figueiredo. An augmented lagrangian  
489 approach to the constrained optimization formulation of imaging inverse problems. *IEEE trans-*  
490 *actions on image processing*, 20(3):681–695, 2010.
- 491 Benjamin Boys, Mark Girolami, Jakiw Pidstrigach, Sebastian Reich, Alan Mosca, and O Deniz  
492 Akyildiz. Tweedie moment projected diffusions for inverse problems. *arXiv preprint*  
493 *arXiv:2310.06721*, 2023.
- 494 Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated mri. *Medical*  
495 *image analysis*, 80:102479, 2022.
- 497 Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion  
498 posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022a.
- 499 Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models  
500 for inverse problems using manifold constraints. *Advances in Neural Information Processing*  
501 *Systems*, 35:25683–25696, 2022b.
- 503 Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating con-  
504 ditional diffusion models for inverse problems through stochastic contraction. In *Proceedings of*  
505 *the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12413–12422, 2022c.
- 506 Hyungjin Chung, Suhyeon Lee, and Jong Chul Ye. Fast diffusion sampler for inverse problems by  
507 geometric decomposition. *arXiv preprint arXiv:2303.05754*, 3(4), 2023.
- 508 Andrew R Conn, Nicholas IM Gould, and Philippe L Toint. *Trust region methods*. SIAM, 2000.
- 509 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-  
510 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,  
511 pp. 248–255. Ieee, 2009.
- 512 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances*  
513 *in neural information processing systems*, 34:8780–8794, 2021.
- 514 Bradley Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Associa-*  
515 *tion*, 106(496):1602–1614, 2011.
- 516 Daniel Gabay and Bertrand Mercier. A dual algorithm for the solution of nonlinear variational  
517 problems via finite element approximation. *Computers & mathematics with applications*, 2(1):  
518 17–40, 1976.
- 519 Yutong He, Naoki Murata, Chieh-Hsin Lai, Yuhta Takida, Toshimitsu Uesaka, Dongjun Kim, Wei-  
520 Hsiang Liao, Yuki Mitsufuji, J Zico Kolter, Ruslan Salakhutdinov, et al. Manifold preserving  
521 guided diffusion. *arXiv preprint arXiv:2311.16424*, 2023.
- 522 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in*  
523 *neural information processing systems*, 33:6840–6851, 2020.
- 524 Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alexandros G Dimakis, and Jon Tamir. Robust  
525 compressed sensing mri with deep generative priors. *Advances in neural information processing*  
526 *systems*, 34:14938–14954, 2021.
- 527 Kyong Hwan Jin, Michael T McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional  
528 neural network for inverse problems in imaging. *IEEE transactions on image processing*, 26(9):  
529 4509–4522, 2017.
- 530 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative  
531 adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
532 *recognition*, pp. 4401–4410, 2019.
- 533 Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration  
534 models. *Advances in neural information processing systems*, 35:23593–23606, 2022.

- 540 Dongjun Kim, Seungjae Shin, Kyungwoo Song, Wanmo Kang, and Il-Chul Moon. Soft truncation:  
541 A universal training technique of score-based diffusion model for high precision score estimation.  
542 *arXiv preprint arXiv:2106.05527*, 2021.
- 543 Kwanyoung Kim and Jong Chul Ye. Noise2score: tweedie’s approach to self-supervised image  
544 denoising without clean images. *Advances in Neural Information Processing Systems*, 34:864–  
545 874, 2021.
- 546 Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Ad-  
547 vances in neural information processing systems*, 34:21696–21707, 2021.
- 548 Henry Li and Marcus Pereira. Solving inverse problems via diffusion optimal control. *Advances in  
549 Neural Information Processing Systems*, 37:73549–73571, 2024.
- 550 Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on  
551 manifolds. *arXiv preprint arXiv:2202.09778*, 2022a.
- 552 Wei Liu, Xin Xia, Lu Xiong, Yishi Lu, Letian Gao, and Zhuoping Yu. Automated vehicle sideslip  
553 angle estimation considering signal measurement characteristic. *IEEE Sensors Journal*, 21(19):  
554 21675–21687, 2021.
- 555 Wei Liu, Karoll Quijano, and Melba M Crawford. Yolov5-tassel: Detecting tassels in rgb uav  
556 imagery with improved yolov5 based on transfer learning. *IEEE Journal of Selected Topics in  
557 Applied Earth Observations and Remote Sensing*, 15:8085–8094, 2022b.
- 558 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast  
559 ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural  
560 information processing systems*, 35:5775–5787, 2022.
- 561 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast  
562 solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research*, pp.  
563 1–22, 2025.
- 564 Chenlin Meng, Robin Rombach, Ruiqi Gao, Diederik Kingma, Stefano Ermon, Jonathan Ho, and  
565 Tim Salimans. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF  
566 conference on computer vision and pattern recognition*, pp. 14297–14306, 2023.
- 567 Xiangming Meng and Yoshiyuki Kabashima. Diffusion model based posterior sampling for noisy  
568 linear inverse problems. *arXiv preprint arXiv:2211.12343*, 2022.
- 569 Taylor R Moen, Baiyu Chen, David R Holmes III, Xinhui Duan, Zhicong Yu, Lifeng Yu, Shuai  
570 Leng, Joel G Fletcher, and Cynthia H McCollough. Low-dose ct image and projection dataset.  
571 *Medical physics*, 48(2):902–911, 2021.
- 572 John L Nazareth. Conjugate-gradient methods. In *Encyclopedia of Optimization*, pp. 466–470.  
573 Springer, 2008.
- 574 Nicki Skafte Detlefsen, Jiri Borovec, Justus Schock, Ananya Harsh, Teddy Koker, Luca Di  
575 Liello, Daniel Stancl, Changsheng Quan, Maxim Grechkin, and William Falcon. TorchMetrics  
576 - Measuring Reproducibility in PyTorch, February 2022. URL [https://github.com/  
577 Lightning-AI/torchmetrics](https://github.com/Lightning-AI/torchmetrics).
- 578 Saiprasad Ravishankar, Jong Chul Ye, and Jeffrey A Fessler. Image reconstruction: From sparsity  
579 to data-adaptive methods and machine learning. *Proceedings of the IEEE*, 108(1):86–109, 2019.
- 580 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
581 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-  
582 ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- 583 Matteo Ronchetti. Torchradon: Fast differentiable routines for computed tomography. *arXiv preprint  
584 arXiv:2009.14788*, 2020.
- 585 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv  
586 preprint arXiv:2202.00512*, 2022.

- 594 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*  
595 *preprint arXiv:2010.02502*, 2020a.
- 596
- 597 Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion  
598 models for inverse problems. In *International Conference on Learning Representations*, 2023.
- 599
- 600 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
601 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint*  
602 *arXiv:2011.13456*, 2020b.
- 603 Paul Suetens. *Fundamentals of medical imaging*. Cambridge university press, 2017.
- 604
- 605 Kaicong Sun and Sven Simon. Bilateral spectrum weighted total variation for noisy-image super-  
606 resolution and image denoising. *IEEE Transactions on Signal Processing*, 69:6329–6341, 2021.
- 607
- 608 Kaicong Sun, Qian Wang, and Dinggang Shen. Joint cross-attention network with deep modality  
609 prior for fast mri reconstruction. *IEEE Transactions on Medical Imaging*, 43(1):558–569, 2023.
- 610
- 611 Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion  
612 null-space model. *arXiv preprint arXiv:2212.00490*, 2022.
- 613
- 614 Lingxiao Yang, Shutong Ding, Yifan Cai, Jingyi Yu, Jingya Wang, and Ye Shi. Guidance with  
615 spherical gaussian constraint for conditional diffusion. In *Proceedings of the 41st International*  
616 *Conference on Machine Learning*, pp. 56071–56095, 2024.
- 617
- 618 Jiwen Yu, Yinhuai Wang, Chen Zhao, Bernard Ghanem, and Jian Zhang. Freedom: Training-free  
619 energy-guided conditional diffusion model. In *Proceedings of the IEEE/CVF International Con-*  
620 *ference on Computer Vision*, pp. 23174–23184, 2023.
- 621
- 622 Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J Muckley,  
623 Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al. fastmri: An open dataset and  
624 benchmarks for accelerated mri. *arXiv preprint arXiv:1811.08839*, 2018.
- 625
- 626 Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song.  
627 Improving diffusion inverse problem solving with decoupled noise annealing. In *Proceedings of*  
628 *the Computer Vision and Pattern Recognition Conference*, pp. 20895–20905, 2025.
- 629
- 630 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable  
631 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on*  
632 *computer vision and pattern recognition*, pp. 586–595, 2018.
- 633
- 634 Kaiwen Zheng, Cheng Lu, Jianfei Chen, and Jun Zhu. Dpm-solver-v3: Improved diffusion ode  
635 solver with empirical model statistics. *Advances in Neural Information Processing Systems*, 36:  
636 55502–55542, 2023.
- 637
- 638
- 639
- 640
- 641
- 642
- 643
- 644
- 645
- 646
- 647

## APPENDIX

## A PROOF FOR PROPOSITIONS

**Proposition 3 (Restated).** For the case of VP-SDE or VE-SDE sampling (Song et al., 2020b),  $q_t(\mathbf{x}_0 | \mathbf{x}_t)$  has the unique posterior mean at

$$\hat{\mathbf{x}}_0(\mathbf{x}_t) = \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \frac{1}{\sqrt{\bar{\alpha}_t}} [\mathbf{x}_t + (1 - \bar{\alpha}_t) \mathbf{s}_\theta(\mathbf{x}_t, t)]. \quad (25)$$

For the case of DDPM sampling (Ho et al., 2020), the formula can be rewritten as

$$\hat{\mathbf{x}}_0(\mathbf{x}_t) = \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \frac{1}{\sqrt{\bar{\alpha}_t}} [\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)]. \quad (26)$$

**Proof.** Consider the forward process  $\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sigma_t z$  with  $z \sim \mathcal{N}(0, I)$ , where  $\sigma_t = \sqrt{1 - \bar{\alpha}_t}$ .

Let  $\mathbf{x} = \sqrt{\bar{\alpha}_t} \mathbf{x}_0$ , we want to compute

$$\mathbb{E}[\mathbf{x} | \mathbf{x}_t] = \int \mathbf{x} q_t(\mathbf{x} | \mathbf{x}_t) d\mathbf{x}. \quad (27)$$

Using Bayes' rule,

$$\mathbb{E}[\mathbf{x} | \mathbf{x}_t] = \int \mathbf{x} \frac{q_t(\mathbf{x}_t | \mathbf{x}) q_t(\mathbf{x})}{q_t(\mathbf{x}_t)} d\mathbf{x} = \frac{1}{q_t(\mathbf{x}_t)} \int \mathbf{x} q_t(\mathbf{x}_t | \mathbf{x}) q_t(\mathbf{x}) d\mathbf{x}. \quad (28)$$

Since  $\mathbf{x}_t | \mathbf{x} \sim \mathcal{N}(\mathbf{x}, \sigma_t^2 I)$ , we have

$$q_t(\mathbf{x}_t | \mathbf{x}) = \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left(-\frac{\|\mathbf{x}_t - \mathbf{x}\|^2}{2\sigma_t^2}\right). \quad (29)$$

Thus,

$$\mathbb{E}[\mathbf{x} | \mathbf{x}_t] = \frac{1}{q_t(\mathbf{x}_t)} \int \mathbf{x} \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left(-\frac{\|\mathbf{x}_t - \mathbf{x}\|^2}{2\sigma_t^2}\right) q_t(\mathbf{x}) d\mathbf{x}. \quad (30)$$

Multiply numerator and denominator by  $\sqrt{\bar{\alpha}_t}$ , and rewrite

$$\mathbf{x} = \sigma_t^2 \frac{\mathbf{x} - \mathbf{x}_t}{\sigma_t^2} + \mathbf{x}_t. \quad (31)$$

So,

$$\mathbb{E}[\mathbf{x} | \mathbf{x}_t] = \frac{1}{q_t(\mathbf{x}_t)} \int \left(\sigma_t^2 \frac{\mathbf{x} - \mathbf{x}_t}{\sigma_t^2} + \mathbf{x}_t\right) \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left(-\frac{\|\mathbf{x}_t - \mathbf{x}\|^2}{2\sigma_t^2}\right) q_t(\mathbf{x}) d\mathbf{x}. \quad (32)$$

This splits into two terms,

$$\begin{aligned} \mathbb{E}[\mathbf{x} | \mathbf{x}_t] &= \frac{1}{q_t(\mathbf{x}_t)} \int \sigma_t^2 \frac{\mathbf{x} - \mathbf{x}_t}{\sigma_t^2} \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left(-\frac{\|\mathbf{x}_t - \mathbf{x}\|^2}{2\sigma_t^2}\right) q_t(\mathbf{x}) d\mathbf{x} \\ &\quad + \frac{1}{q_t(\mathbf{x}_t)} \int \mathbf{x}_t \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left(-\frac{\|\mathbf{x}_t - \mathbf{x}\|^2}{2\sigma_t^2}\right) q_t(\mathbf{x}) d\mathbf{x}. \end{aligned} \quad (33)$$

Now, notice that the first term is

$$\begin{aligned} &\frac{1}{q_t(\mathbf{x}_t)} \int \sigma_t^2 \frac{\mathbf{x} - \mathbf{x}_t}{\sigma_t^2} \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left(-\frac{\|\mathbf{x}_t - \mathbf{x}\|^2}{2\sigma_t^2}\right) q_t(\mathbf{x}) d\mathbf{x} \\ &= \frac{\sigma_t^2}{q_t(\mathbf{x}_t)} \int \frac{d\left[\frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left(-\frac{\|\mathbf{x}_t - \mathbf{x}\|^2}{2\sigma_t^2}\right)\right]}{d\mathbf{x}_t} q_t(\mathbf{x}) d\mathbf{x} \\ &= \frac{\sigma_t^2}{q_t(\mathbf{x}_t)} \frac{d}{d\mathbf{x}_t} \int q_t(\mathbf{x}_t | \mathbf{x}) q_t(\mathbf{x}) d\mathbf{x} \\ &= \frac{\sigma_t^2}{q_t(\mathbf{x}_t)} \frac{dq_t(\mathbf{x}_t)}{d\mathbf{x}_t} \\ &= \sigma_t^2 \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t), \end{aligned} \quad (34)$$

702 and the second term is

$$\begin{aligned}
703 & \\
704 & \frac{1}{q_t(\mathbf{x}_t)} \int \mathbf{x}_t \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left(-\frac{\|\mathbf{x}_t - \mathbf{x}\|^2}{2\sigma_t^2}\right) q_t(\mathbf{x}) d\mathbf{x} \\
705 & \\
706 & = \frac{1}{q_t(\mathbf{x}_t)} \int \mathbf{x}_t q_t(\mathbf{x}_t | \mathbf{x}) q_t(\mathbf{x}) d\mathbf{x} \\
707 & \\
708 & = \mathbf{x}_t \\
709 &
\end{aligned} \tag{35}$$

710 Hence,

$$711 \mathbb{E}[\mathbf{x} | \mathbf{x}_t] = \mathbf{x}_t + \sigma_t^2 \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t) \tag{36}$$

712 Finally, we expand  $\mathbf{x} = \sqrt{\bar{\alpha}_t} \mathbf{x}_0$  and invoke the linearity of the expectation to arrive at

$$713 \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \frac{1}{\sqrt{\bar{\alpha}_t}} [\mathbf{x}_t + \sigma_t^2 \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t)] \tag{37}$$

714 In diffusion models, the  $\epsilon$ -prediction parameterization estimates the injected noise term  $\epsilon_\theta$ , whereas the score-prediction parameterization estimates the score function  $\mathbf{s}_\theta$ . Their relationship is given by

$$\begin{aligned}
715 & \epsilon_\theta = -\sigma_t \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t), \\
716 & \mathbf{s}_\theta = \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t)
\end{aligned} \tag{38}$$

717 Accordingly, the posterior mean of the clean sample  $\mathbf{x}_0$  given  $\mathbf{x}_t$  can be expressed in two equivalent forms:

$$\begin{aligned}
718 & \hat{\mathbf{x}}_0(\mathbf{x}_t) = \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \frac{1}{\sqrt{\bar{\alpha}_t}} [\mathbf{x}_t + (1 - \bar{\alpha}_t) \mathbf{s}_\theta(\mathbf{x}_t, t)], \\
719 & \\
720 & \hat{\mathbf{x}}_0(\mathbf{x}_t) = \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \frac{1}{\sqrt{\bar{\alpha}_t}} [\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t)].
\end{aligned}$$

721 □

722 **Proposition 4 (Restated).** *In conditional diffusion sampling, once the estimate  $\hat{\mathbf{x}}_0$  incorporates measurement information  $\mathbf{y}$ , the posterior distribution  $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \hat{\mathbf{x}}_0)$  can be well approximated by the marginal distribution  $q(\mathbf{x}_{t-1} | \hat{\mathbf{x}}_0)$ .*

723 **Proof.** In unconditional diffusion sampling, substituting the clean signal  $\mathbf{x}_0$  into the posterior distribution

$$724 q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0), \tilde{\beta}_t I), \quad \tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t, \quad \beta_t = 1 - \alpha_t \tag{39}$$

725 and the true reverse posterior mean is

$$726 \tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0) = \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_t}{1 - \bar{\alpha}_t} \mathbf{x}_0 + \frac{\sqrt{\bar{\alpha}_t} (1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \mathbf{x}_t. \tag{40}$$

727 Since  $\mathbf{x}_0$  and  $\mathbf{x}_t$  come from the same forward process, they are matched. Integrating over  $\mathbf{x}_0$  gives the exact marginal

$$728 q_{t-1}(\mathbf{x}_{t-1}) \propto \int q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) q(\mathbf{x}_0 | \mathbf{x}_t) d\mathbf{x}_0. \tag{41}$$

729 Thus, the marginal distribution  $q_{t-1}(\mathbf{x}_{t-1})$  is exact in the unconditional case.

730 However, in conditional sampling with measurement guidance, when  $\mathbf{x}_0$  is replaced by the estimator  $\hat{\mathbf{x}}_0(\mathbf{x}_t, \mathbf{y})$ , the update

$$731 q(\mathbf{x}_{t-1} | \mathbf{x}_t, \hat{\mathbf{x}}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\mu}_t(\mathbf{x}_t, \hat{\mathbf{x}}_0), \tilde{\beta}_t I). \tag{42}$$

732 Substituting into the bridge posterior yields

$$733 \tilde{\mu}_t(\mathbf{x}_t, \hat{\mathbf{x}}_0) = \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_t}{1 - \bar{\alpha}_t} \hat{\mathbf{x}}_0 + \frac{\sqrt{\bar{\alpha}_t} (1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \mathbf{x}_t. \tag{43}$$

Here  $\hat{x}_0(\mathbf{x}_t, \mathbf{y})$  encodes information from  $\mathbf{y}$ , but the noise term  $\epsilon_\theta(\mathbf{x}_t, t)$  (hidden inside the dependence on  $\mathbf{x}_t$ ) is still unconditional. This mismatch introduces an error in the mean:

$$\Delta_t = \tilde{\mu}_t(\mathbf{x}_t, \hat{x}_0) - \tilde{\mu}_t(\mathbf{x}_t, x_0). \tag{44}$$

Due to  $\hat{x}_0(\mathbf{x}_t, \mathbf{y}) \approx \arg \min_{\hat{x}_0} \|\mathbf{y} - \mathcal{A}(\hat{x}_0)\|^2$ , and  $\hat{x}_0$  is independent of  $\mathbf{x}_t$ , the marginal distribution  $q_{t-1}(\mathbf{x}_{t-1})$  approximation discards the dependence on  $\mathbf{x}_t$

$$q_{t-1}(\mathbf{x}_{t-1}) \propto q(\mathbf{x}_{t-1} | \hat{x}_0) q(\hat{x}_0), \tag{45}$$

which depends solely on the quality of  $\hat{x}_0$  and avoids the inconsistency term from  $\mathbf{x}_t$ . Consequently,  $q(\mathbf{x}_{t-1} | \hat{x}_0)$  provides a cleaner and measurement-consistent reverse step.

□

## B ADDITIONAL EXPERIMENTS

### B.1 EXPERIMENTAL RESULTS ON FASTMRI AND AAPM LDCT

Figure 4, 5, 6, provide additional visualization results for Brain MRI, Knee MRI, and CT reconstruction tasks under different undersampling settings. These examples complement the quantitative evaluations in the main text and highlight the ability of CGRS to produce high-fidelity reconstructions with sharp structural details and reduced artifacts across diverse tasks.

### B.2 EXPERIMENTAL RESULTS ON FFHQ AND IMAGENET

Method	PSNR↑	SSIM↑	LPIPS↓
Diff-OC	25.24	0.738	0.374
DDRM	24.43	0.721	0.393
DPS	27.50	0.797	0.352
DDNM	27.68	0.815	0.364
DAPS	<u>29.94</u>	<u>0.827</u>	<u>0.355</u>
CGRS	<b>35.79</b>	<b>0.953</b>	<b>0.317</b>

Table 5: **Quantitative evaluation of Gaussian deblurring on FFHQ.** The best and second-best results within each type of task are indicated by bold and underlined marks, respectively.

Method	Super-Resolution (×4)			Gaussian Deblurring		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Diff-OC	22.28	<u>0.615</u>	0.398	22.19	0.539	0.425
DDRM	20.12	0.586	0.443	20.07	0.522	0.451
DPS	23.05	0.612	<u>0.369</u>	22.19	0.539	<u>0.396</u>
DDNM	21.54	0.597	0.436	21.37	0.527	0.446
DAPS	<u>23.22</u>	0.594	0.396	<u>22.94</u>	<u>0.574</u>	0.413
CGRS	<b>23.88</b>	<b>0.663</b>	<b>0.354</b>	<b>25.47</b>	<b>0.757</b>	<b>0.357</b>

Table 6: **Quantitative evaluation on ImageNet.** The best and second-best results within each type of task are indicated by bold and underlined marks, respectively.

Tab 5, 6 provide additional quantitative results on FFHQ and ImageNet datasets, respectively. CGRS outperforms existing diffusion-based baselines. Figure 7, 8, 9, 10 provide more visualization results for super-resolution and Gaussian deblurring tasks.

810

811

812

813

814

815

816

817

818

819

820

821

822

823

824

825

826

827

828

829

830

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

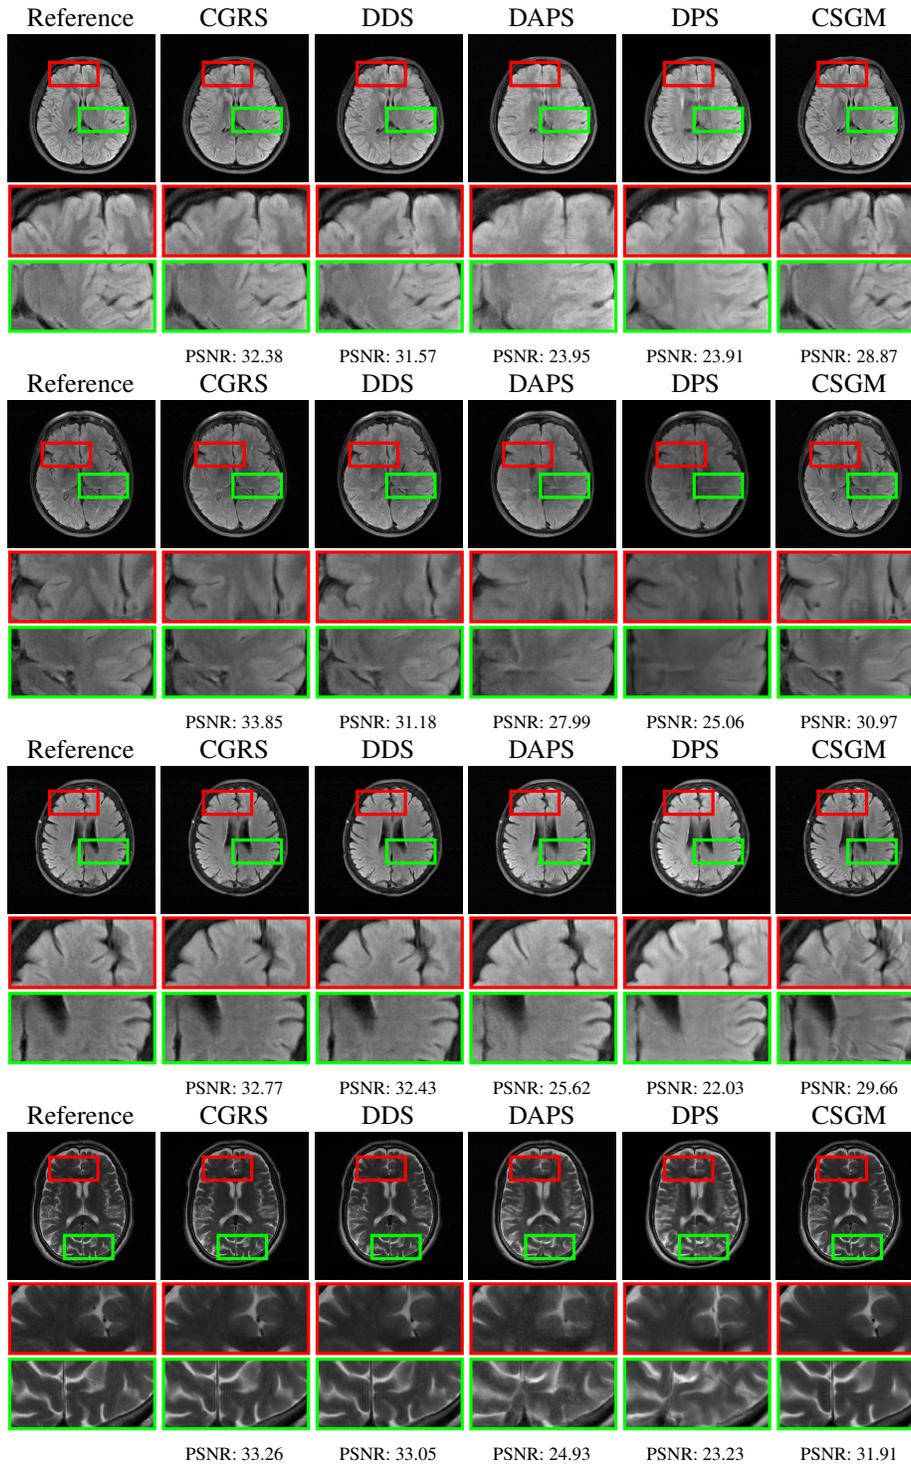


Figure 4: More samples for Brain MRI reconstruction under  $\times 8$  acceleration. CGRS is able to generate high-fidelity reconstructions for MRI tasks.

864

865

866

867

868

869

870

871

872

873

874

875

876

877

878

879

880

881

882

883

884

885

886

887

888

889

890

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

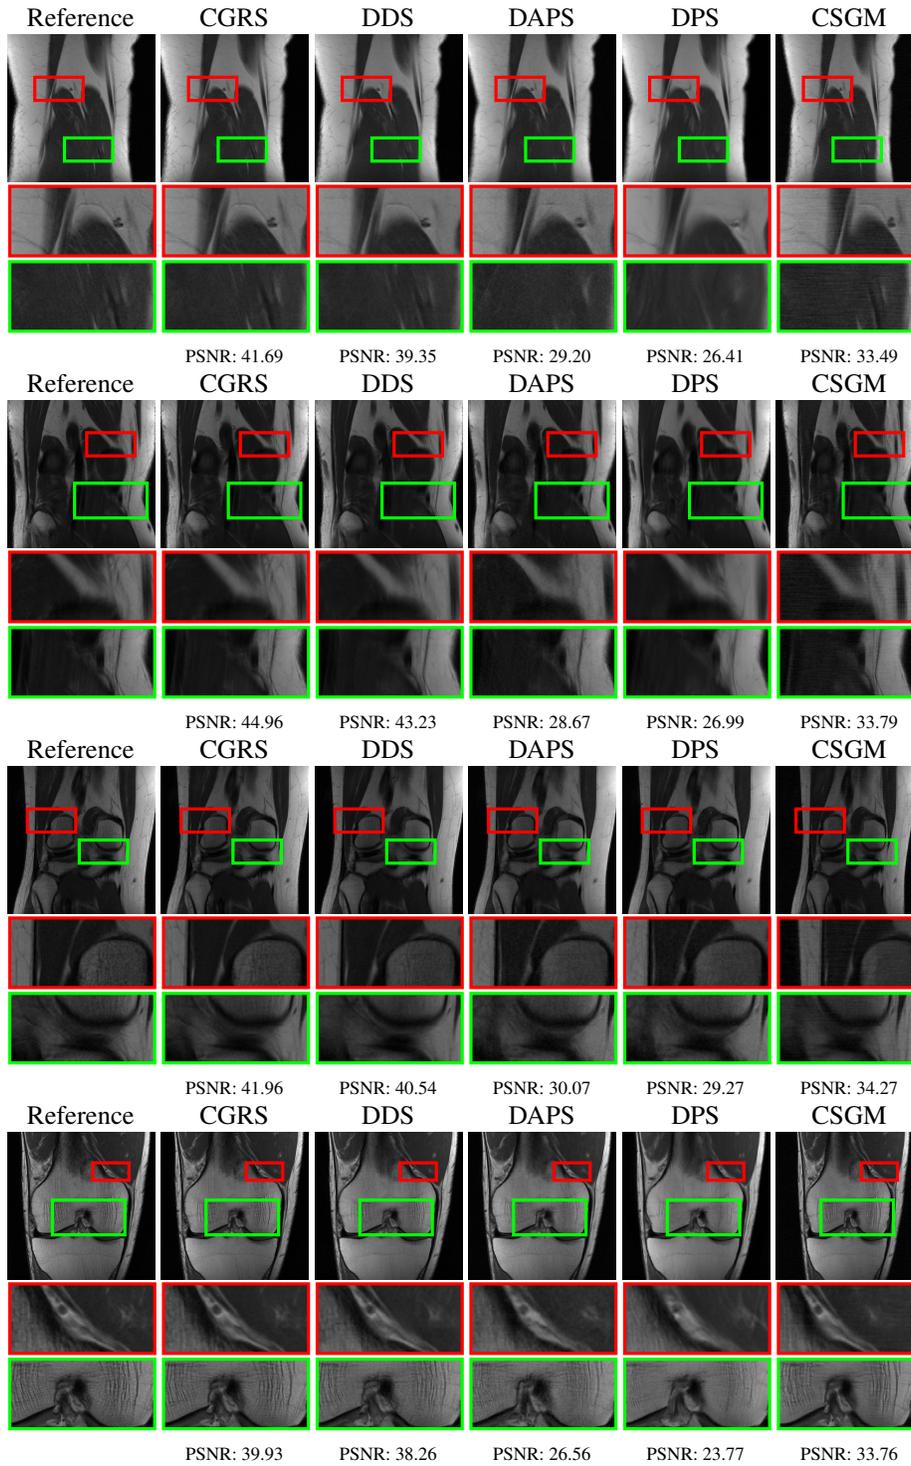


Figure 5: More samples for Knee MRI reconstruction under  $\times 4$  acceleration. CGRS is able to generate high-fidelity reconstructions for MRI tasks.

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

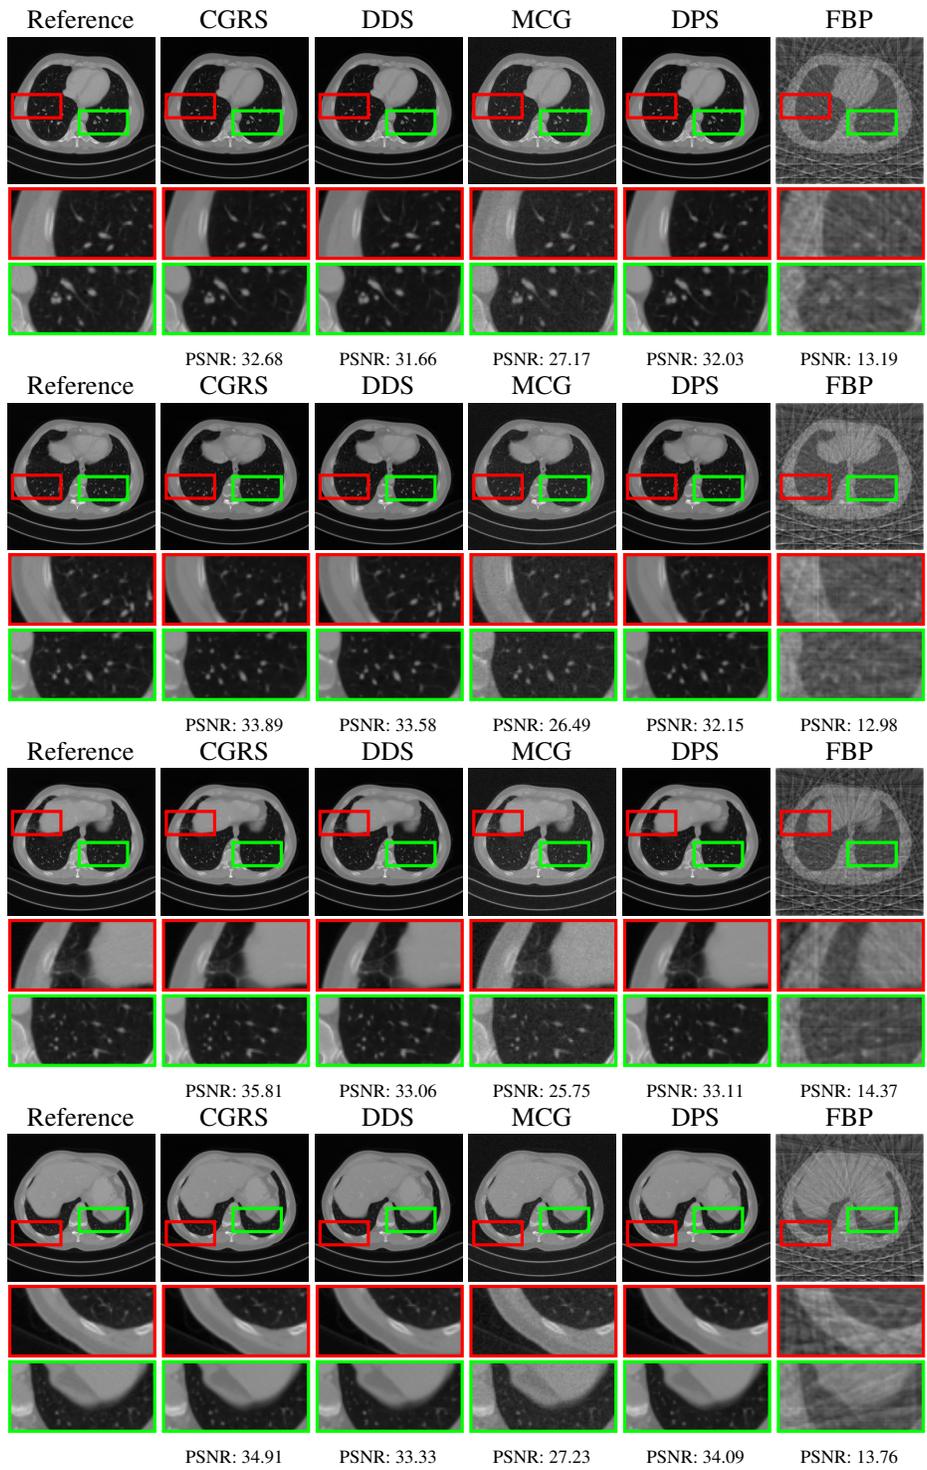


Figure 6: More samples for CT reconstruction using 18 projection angles.

972

973

974

975

976

977

978

979

980

981

982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

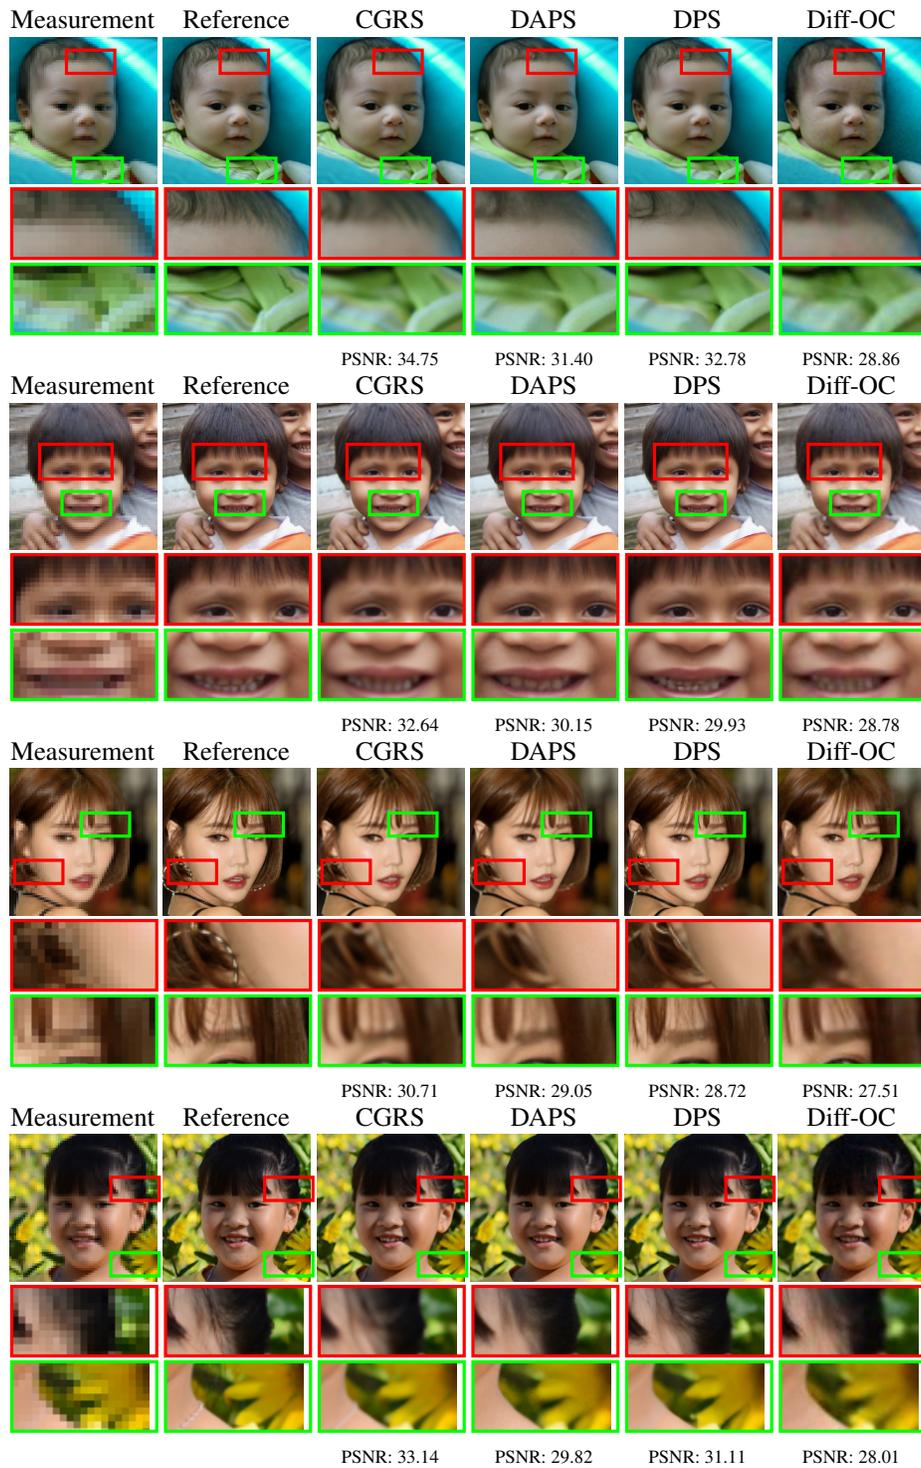


Figure 7: Qualitative results of super-resolution  $\times 4$  on FFHQ.

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

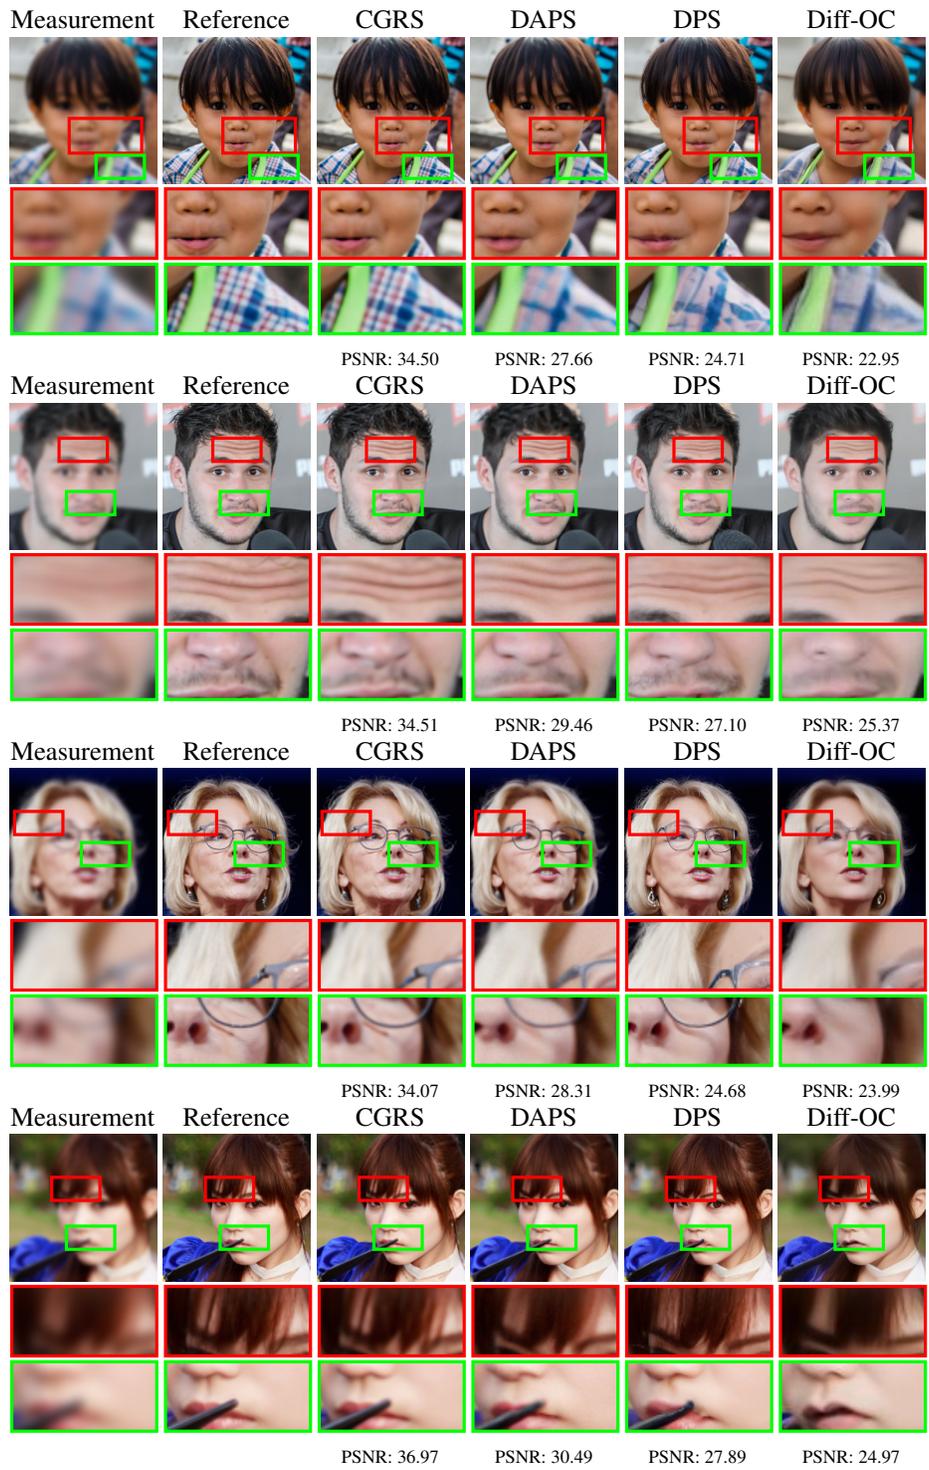


Figure 8: Qualitative results of Gaussian deblurring on FFHQ.

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

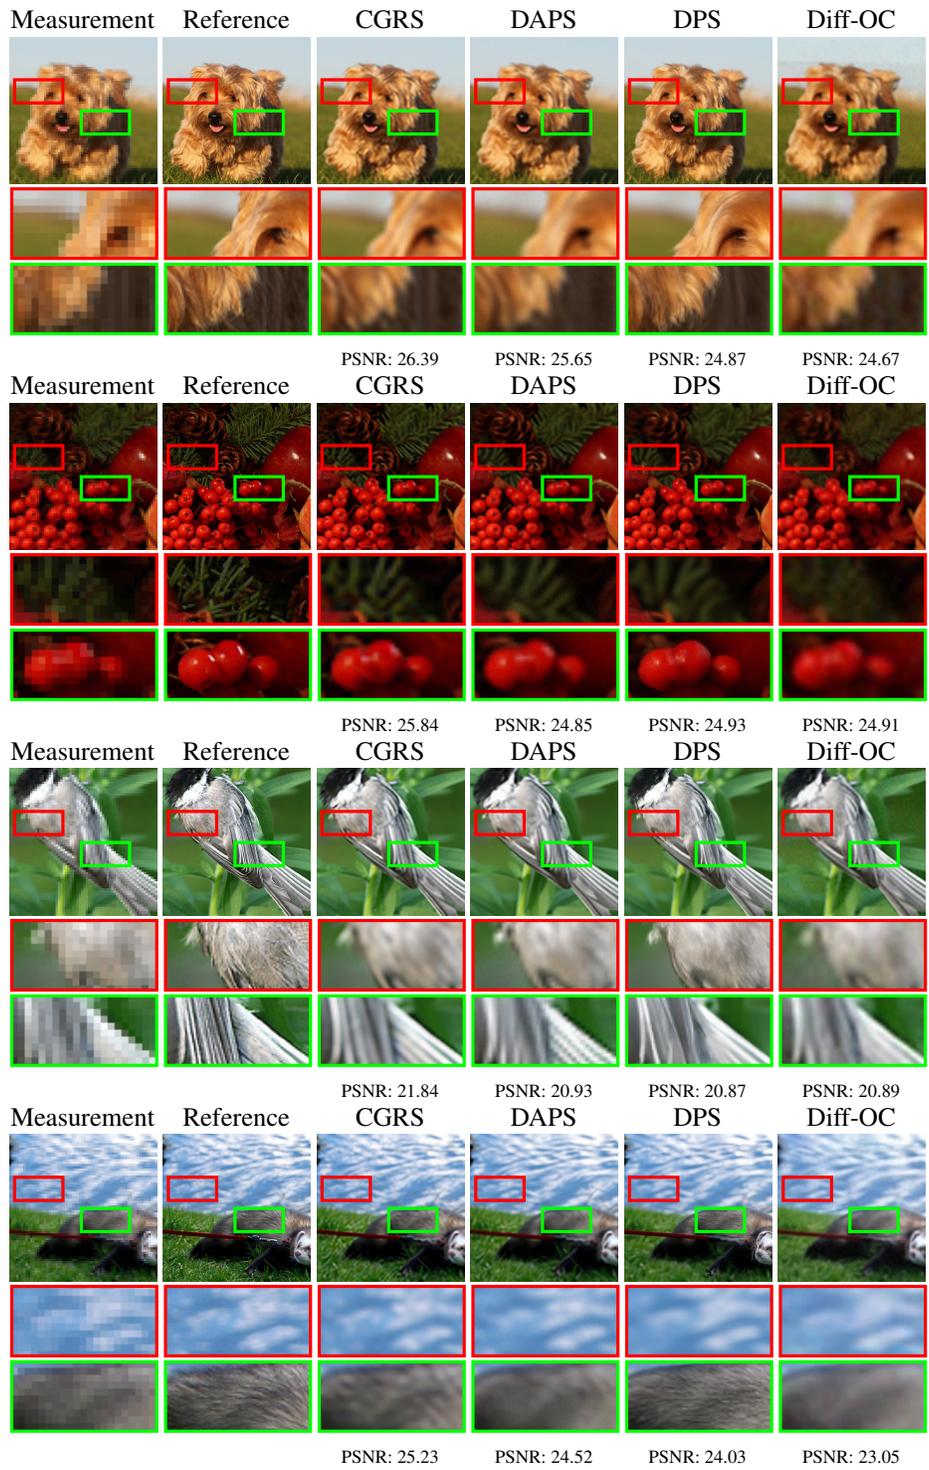


Figure 9: Qualitative results of super-resolution  $\times 4$  on ImageNet.

1134  
 1135  
 1136  
 1137  
 1138  
 1139  
 1140  
 1141  
 1142  
 1143  
 1144  
 1145  
 1146  
 1147  
 1148  
 1149  
 1150  
 1151  
 1152  
 1153  
 1154  
 1155  
 1156  
 1157  
 1158  
 1159  
 1160  
 1161  
 1162  
 1163  
 1164  
 1165  
 1166  
 1167  
 1168  
 1169  
 1170  
 1171  
 1172  
 1173  
 1174  
 1175  
 1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187

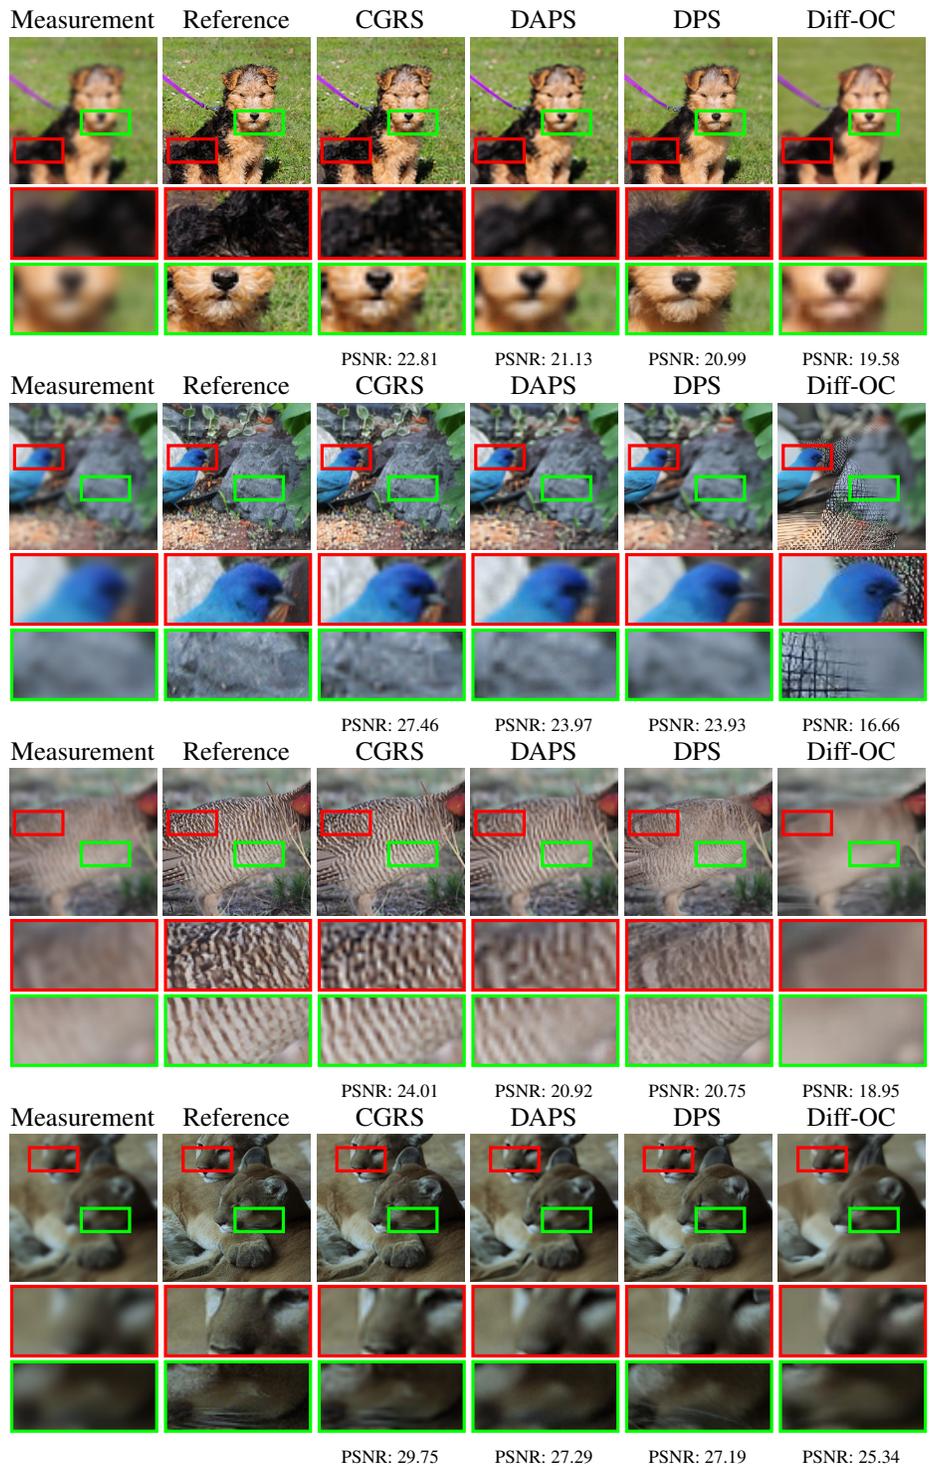


Figure 10: Qualitative results of Gaussian deblurring on ImageNet.