

Solving Inverse Problems using Diffusion with Iterative Colored Renoising

Anonymous authors

Paper under double-blind review

Abstract

Imaging inverse problems can be solved in an unsupervised manner using pre-trained diffusion models, but doing so requires approximating the gradient of the measurement-conditional score function in the diffusion reverse process. We show that the approximations produced by existing methods are relatively poor, especially early in the reverse process, and so we propose a new approach that iteratively reestimates and “renoises” the estimate several times per diffusion step. This iterative approach, which we call Fast Iterative REnoising (FIRE), injects colored noise that is shaped to ensure that the pre-trained diffusion model always sees white noise, in accordance with how it was trained. We then embed FIRE into the DDIM reverse process and show that the resulting “DDfire” offers state-of-the-art accuracy and runtime on several linear inverse problems, as well as phase retrieval.

1 Introduction

Diffusion has emerged as a powerful approach to generate samples from a complex distribution p_0 (Sohl-Dickstein et al., 2015; Song & Ermon, 2019; Ho et al., 2020; Song et al., 2021b;a). Recently, diffusion has also been used to solve inverse problems (Daras et al., 2024), where the goal is to recover $\mathbf{x}_0 \sim p_0$ from incomplete, distorted, and/or noisy measurements \mathbf{y} in an unsupervised manner. There, a diffusion model is trained to generate samples from p_0 and, at test time, the reverse process is modified to incorporate knowledge of the measurements \mathbf{y} , with the goal of sampling from the posterior distribution $p(\mathbf{x}_0|\mathbf{y})$.

When implementing the reverse process, the main challenge is approximating the conditional score function $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y})$ at each step t , where \mathbf{x}_t is a additive-white-Gaussian-noise (AWGN) corrupted and possibly scaled version of $\mathbf{x}_0 \in \mathbb{R}^d$, and $\mathbf{y} \in \mathbb{R}^m$ is treated as a draw from a likelihood function $p(\mathbf{y}|\mathbf{x}_0)$. (See Sec. 2 for more details.) Many existing approaches fall into one of two categories. The first uses Bayes rule to write $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y}) = \nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}} \ln p_t(\mathbf{y}|\mathbf{x}_t)$, where an approximation of $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t)$ is readily available from the p_0 -trained diffusion model, and then approximates $\nabla_{\mathbf{x}} \ln p_t(\mathbf{y}|\mathbf{x}_t)$ (e.g., Chung et al. (2023a); Song et al. (2023)). The second approach uses Tweedie’s formula (Efron, 2011) to write

$$\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y}) = \frac{\mathbb{E}\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\} - \mathbf{x}_t}{\sigma_t^2} \quad (1)$$

and then approximates the conditional denoiser $\mathbb{E}\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\}$ (e.g., Kavar et al. (2022a); Wang et al. (2023); Zhu et al. (2023); Chung et al. (2024)).

A key shortcoming of the aforementioned approaches is that their conditional-score approximations are not very accurate, especially early in the reverse process. For the methods that approximate $\mathbb{E}\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\}$, we can assess the approximation quality both visually and via mean-square error (MSE) or PSNR, since the exact $\mathbb{E}\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\}$ minimizes MSE given \mathbf{x}_t and \mathbf{y} . For the methods that approximate $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y})$, we can compute their equivalent conditional-denoiser approximations using

$$\mathbb{E}\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\} = \mathbf{x}_t + \sigma_t^2 \nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y}), \quad (2)$$

which follows from (1). Figure 1 shows $\mathbb{E}\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\}$ -approximations from the DDRM (Kavar et al., 2022a), DiffPIR (Zhu et al., 2023), DPS (Chung et al., 2023a), and DAPS Zhang et al. (2025) solvers at times 25%,

50%, and 75% through their reverse processes for noisy box inpainting with $\sigma_y = 0.05$. The approximations show unwanted artifacts, especially early in the reverse process.




Figure 1: Left column: True x_0 , noisy box inpainting y , and 50-iteration FIRE approximation of $E\{x_0|y\}$. Other columns: Approximations of $E\{x_0|x_t, y\}$ at different t (as measured by % NFEs). Note the over-smoothing with DDRM and DPS. Additionally, note the cut-and-paste artifacts of DiffPIR and DAPS.

In this paper, we aim to improve the approximation of $E\{x_0|x_t, y\}$ at each step t . Since we have observed that methods using a single neural function evaluation (NFE) to approximate $E\{x_0|x_t, y\}$ or $\nabla_x \ln p_t(x_t|y)$ perform poorly, we consider using several NFEs. In particular, we propose an iterative approach to approximating $E\{x_0|x_t, y\}$ that we call Fast Iterative REnoising (FIRE). FIRE is like a plug-and-play (PnP) algorithm (see the PnP survey Ahmad et al. (2020)) in that it iterates unconditional denoising with linear estimation from x_t and y . But unlike traditional PnP algorithms, which aim to minimize an implicit loss function, FIRE is based on *colored renoising*, where carefully designed colored Gaussian noise is added to the linear-estimation output so that the denoiser’s input error approximates AWGN. Since the denoiser is trained to remove AWGN, renoising aims to minimize the distribution shift experienced during inference. Figure 1 shows the 50-iteration FIRE approximation to $E\{x_0|x_T, y\} = E\{x_0|y\}$ for noisy box inpainting, which is relatively free of artifacts. We propose two versions of FIRE: one that handles AWGN-corrupted linear measurement models like (6), and one that handles generalized-linear inverse problems such as phase retrieval (Shechtman et al., 2015), Poisson regression (Figueiredo & Bioucas-Dias, 2010), and dequantization (Zymnis et al., 2010). The latter is based on expectation propagation (EP) (Minka, 2001; Bishop, 2007).

We then embed FIRE into the DDIM diffusion reverse process (Song et al., 2021a), yielding the “DDfire” posterior sampler. Figure 1 shows examples of DDfire’s $E\{x_0|x_t, y\}$ approximation when run on noisy box inpainting. DDfire’s $E\{x_0|x_t, y\}$ approximations have fewer structural artifacts and higher PSNR than its competitors.

The contributions of this work are as follows:

1. We propose linear FIRE, an iterative approach to solving AWGN-corrupted linear inverse problems that injects carefully designed colored Gaussian noise in order to whiten the denoiser’s input error.
2. We theoretically analyze the convergence of linear FIRE.
3. We use expectation propagation (EP) to extend linear FIRE to generalized-linear inverse problems.
4. Combining FIRE and DDIM, we propose the DDfire diffusion posterior sampler.
5. We demonstrate the excellent accuracy of DDfire on several imaging inverse problems: box inpainting, Gaussian and motion blur, super-resolution, and phase retrieval.

2 Background

Given training data drawn from distribution p_0 , diffusion models corrupt the data with ever-increasing amounts of noise and then learn to reverse that process in a way that can generate new samples from p_0 . In the main text, we assume the variance-exploding (VE) diffusion formulation (Song et al., 2021b), whereas Appendix A provides details on the variance-preserving (VP) formulation, including DDPM (Ho et al., 2020) and DDIM (Song et al., 2021a).

The VE diffusion forward process can be written as a stochastic differential equation (SDE) $d\mathbf{x} = \sqrt{d[\sigma^2(t)]}/dt d\mathbf{w}$ over t from 0 to T , where $\sigma^2(t)$ is a variance schedule and $d\mathbf{w}$ is the standard Wiener process (SWP) (Song et al., 2021b). The corresponding reverse process runs the SDE $d\mathbf{x} = -\sigma^2(t)\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}) dt + \sqrt{d[\sigma^2(t)]}/dt d\bar{\mathbf{w}}$ backwards over t from T to 0, where $p_t(\cdot)$ is the marginal distribution of \mathbf{x} at t and $d\bar{\mathbf{w}}$ is the SWP run backwards. The “score function” $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x})$ can be approximated using a deep neural network (DNN) $\mathbf{s}_{\theta}(\mathbf{x}, t)$ trained via denoising score matching (Hyvärinen, 2005).

In practice, time is discretized to $t \in \{0, 1, \dots, T\}$, yielding the SMLD from Song & Ermon (2019), whose forward process, $\mathbf{x}_{t+1} = \mathbf{x}_t + \sqrt{\sigma_{t+1}^2 - \sigma_t^2} \mathbf{w}_t$, with i.i.d $\{\mathbf{w}_t\} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\sigma_0^2 = 0$, implies that

$$\mathbf{x}_t = \mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (3)$$

for all $t \in \{0, 1, \dots, T\}$. The SMLD reverse process then uses i.i.d $\{\mathbf{n}_t\} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ in

$$\mathbf{x}_t = \mathbf{x}_{t+1} + (\sigma_{t+1}^2 - \sigma_t^2) \nabla_{\mathbf{x}} \ln p_{t+1}(\mathbf{x}_{t+1}) + \sqrt{\frac{\sigma_t^2(\sigma_{t+1}^2 - \sigma_t^2)}{\sigma_{t+1}^2}} \mathbf{n}_{t+1}. \quad (4)$$

To exploit side information about \mathbf{x}_0 , such as the measurements \mathbf{y} in an inverse problem, one can simply replace $p_t(\cdot)$ with $p_t(\cdot|\mathbf{y})$ in the above equations (Song et al., 2021b). However, most works aim to avoid training a \mathbf{y} -dependent approximation of the conditional score function $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y})$. Rather, they take an “unsupervised” approach, where $\mathbf{s}_{\theta}(\mathbf{x}_t, t) \approx \nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t)$ is learned during training but \mathbf{y} is presented only at inference (Daras et al., 2024). In this case, approximating $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y})$ is the key technical challenge.

There are two major approaches to approximate $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y})$. The first uses the Bayes rule to write $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y}) = \nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}} \ln p_t(\mathbf{y}|\mathbf{x}_t)$ and then replaces $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t)$ with the score approximation $\mathbf{s}_{\theta}(\mathbf{x}_t, t)$. But the remaining term, $\nabla_{\mathbf{x}} \ln p_t(\mathbf{y}|\mathbf{x}_t)$, is intractable because $p_t(\mathbf{y}|\mathbf{x}_t) = \int p(\mathbf{y}|\mathbf{x}_0)p(\mathbf{x}_0|\mathbf{x}_t) d\mathbf{x}_0$ with unknown $p(\mathbf{x}_0|\mathbf{x}_t)$, and so several approximations have been proposed. For example, DPS (Chung et al., 2023a) uses $p(\mathbf{x}_0|\mathbf{x}_t) \approx \delta(\mathbf{x}_0 - \hat{\mathbf{x}}_{0|t})$, where $\hat{\mathbf{x}}_{0|t}$ is the approximation of $E\{\mathbf{x}_0|\mathbf{x}_t\}$ computed from $\mathbf{s}_{\theta}(\mathbf{x}_t, t)$ using Tweedie’s formula:

$$\hat{\mathbf{x}}_{0|t} = \mathbf{x}_t + \sigma_t^2 \mathbf{s}_{\theta}(\mathbf{x}_t, t). \quad (5)$$

Similarly, IIGDM (Song et al., 2023) uses $p(\mathbf{x}_0|\mathbf{x}_t) \approx \mathcal{N}(\mathbf{x}_0; \hat{\mathbf{x}}_{0|t}, \zeta_t \mathbf{I})$ with some ζ_t . However, a drawback to both approaches is that they require backpropagation through $\mathbf{s}_{\theta}(\cdot, t)$, which increases the cost of generating a single sample. In Fig. 4, we show that DDfire offers a $1.5\times$ speedup over DPS at an equal number of NFEs.

The second major approach to approximating $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y})$ uses (1) with $E\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\}$ approximated by a quantity that we’ll refer to as $\hat{\mathbf{x}}_{0|t, \mathbf{y}}$. For example, with AWGN-corrupted linear measurements

$$\mathbf{y} = \mathbf{A}\mathbf{x}_0 + \sigma_y \mathbf{w} \in \mathbb{R}^m, \quad \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (6)$$

DDNM (Wang et al., 2023) approximates $E\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\}$ by first computing $\hat{\mathbf{x}}_{0|t}$ from (5) and then performing the hard data-consistency step $\hat{\mathbf{x}}_{0|t, \mathbf{y}} = \mathbf{A}^+ \mathbf{y} + (\mathbf{I} - \mathbf{A}^+ \mathbf{A}) \hat{\mathbf{x}}_{0|t}$, where $(\cdot)^+$ is the pseudo-inverse. DDS (Chung et al., 2024) and DiffPIR (Zhu et al., 2023) instead use the soft data-consistency step $\hat{\mathbf{x}}_{0|t, \mathbf{y}} = \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 + \gamma_t \|\mathbf{x} - \hat{\mathbf{x}}_{0|t}\|^2$ with some $\gamma_t > 0$. DDRM (Kawar et al., 2022a) is a related technique that requires a singular value decomposition (SVD), which is prohibitive in many applications.

There are, of course, many other ways to design diffusion posterior samplers, as detailed in the recent overview Daras et al. (2024). For example, RED-diff (Mardani et al., 2024) uses a stochastic version of

the RED algorithm (Romano et al., 2017), whose regularization is the score function (Reehorst & Schniter, 2019). Another approach is to use Markov-chain Monte Carlo (MCMC) methods as inner iterations of the reverse process (Coeurdoux et al., 2024; Wu et al., 2024; Xu & Chi, 2024; Zhang et al., 2025).

3 Approach

As described in Sec. 1, we aim to accurately approximate the conditional denoiser $E\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\}$ at each step of the diffusion reverse process. We first treat the standard linear model (SLM) (6) and later (in Sec. 3.2) consider the generalized linear model (GLM).

3.1 Fast iterative renoising (FIRE) for the SLM

In this section, we describe the FIRE algorithm, which approximates $E\{\mathbf{x}_0|\mathbf{r}_{\text{init}}, \mathbf{y}\}$ assuming \mathbf{y} from (6) and $\mathbf{r}_{\text{init}} = \mathbf{x}_0 + \sigma_{\text{init}}\boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and some $\sigma_{\text{init}} > 0$. The basic FIRE algorithm iterates the following steps $N \geq 1$ times, after initializing $\mathbf{r} \leftarrow \mathbf{r}_{\text{init}}$ and $\sigma \leftarrow \sigma_{\text{init}}$:

- S1) Denoise \mathbf{r} assuming AWGN of variance σ^2 , giving $\bar{\mathbf{x}}$.
- S2) MMSE estimate \mathbf{x}_0 given \mathbf{y} from (6) and the Gaussian prior model $\mathbf{x}_0 \sim \mathcal{N}(\bar{\mathbf{x}}, \nu\mathbf{I})$, giving $\hat{\mathbf{x}}$.
- S3) Update the noise variance via $\sigma^2 \leftarrow \sigma^2/\rho$ with some $\rho > 1$.
- S4) Update $\mathbf{r} \leftarrow \hat{\mathbf{x}} + \mathbf{c}$ using colored Gaussian noise \mathbf{c} created to ensure $\text{Cov}\{\mathbf{r} - \mathbf{x}_0\} = \sigma^2\mathbf{I}$.

Iterating S1)-S3) alone would constitute half-quadratic splitting (HQS) plug-and-play, similar to Zhang et al. (2021). We propose S4), which ensures that the error in \mathbf{r} is white. This is important because denoisers are typically trained to remove white noise, whereas the error in $\hat{\mathbf{x}}$ can be strongly colored. In Sec. 4.1, we demonstrate the advantage of S4). Next we provide details and enhancements of the basic algorithm.

In the sequel, we use “ $\mathbf{d}_\theta(\mathbf{x}, \sigma)$ ” to denote a neural-net approximation of the conditional-mean denoiser $E\{\mathbf{x}_0|\mathbf{x}\}$ of $\mathbf{x} = \mathbf{x}_0 + \sigma\boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Given a score function approximation $\mathbf{s}_\theta(\mathbf{x}, t) \approx \nabla_{\mathbf{x}} \ln p_t(\mathbf{x})$ as discussed in Sec. 2, the denoiser can be constructed via (5) as

$$\mathbf{d}_\theta(\mathbf{x}, \sigma) = \mathbf{x} + \sigma^2 \mathbf{s}_\theta(\mathbf{x}, t) \text{ with } t \text{ such that } \sigma_t = \sigma. \quad (7)$$

When FIRE estimates \mathbf{x}_0 from the measurements \mathbf{y} and the denoiser output $\bar{\mathbf{x}}$, it employs a Gaussian approximation of the form $\mathbf{x}_0 \sim \mathcal{N}(\bar{\mathbf{x}}, \nu\mathbf{I})$, similar to DDS (Chung et al., 2024) and DiffPIR (Zhu et al., 2023). But it differs from DDS and DiffPIR in that ν is explicitly estimated. Note that the Gaussian approximation $\mathbf{x}_0 \sim \mathcal{N}(\bar{\mathbf{x}}, \nu\mathbf{I})$ is equivalent to assuming that

$$\mathbf{x}_0 = \bar{\mathbf{x}} + \sqrt{\nu}\mathbf{e}, \quad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (8)$$

Suppose $\mathbf{x}_0 = \bar{\mathbf{x}} + \sqrt{\nu_0}\mathbf{e}$ with $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, where ν_0 denotes the true error variance. Then (6) and (8) imply

$$E\{\|\mathbf{y} - \mathbf{A}\bar{\mathbf{x}}\|^2\} = E\{\|\mathbf{A}\mathbf{x}_0 + \sigma_y\mathbf{w} - \mathbf{A}\mathbf{x}_0 + \sqrt{\nu_0}\mathbf{A}\mathbf{e}\|^2\} = E\{\|\sigma_y\mathbf{w} + \sqrt{\nu_0}\mathbf{A}\mathbf{e}\|^2\} = m\sigma_y^2 + \nu_0\|\mathbf{A}\|_F^2, \quad (9)$$

assuming independence between \mathbf{e} and \mathbf{w} . Consequently, an unbiased estimate of ν_0 can be constructed as

$$(\|\mathbf{y} - \mathbf{A}\bar{\mathbf{x}}\|^2 - m\sigma_y^2)/\|\mathbf{A}\|_F^2 \triangleq \nu. \quad (10)$$

Figure 8 shows that, in practice, the estimate (10) accurately tracks the true error variance $\|\mathbf{x}_0 - \bar{\mathbf{x}}\|^2/d$.

Under (6) and (8), the MMSE estimate of \mathbf{x}_0 from \mathbf{y} and $\bar{\mathbf{x}}$ can be written as (Poor, 1994)

$$\hat{\mathbf{x}} \triangleq \arg \min_{\mathbf{x}} \left\{ \frac{1}{2\sigma_y^2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 + \frac{1}{2\nu} \|\mathbf{x} - \bar{\mathbf{x}}\|^2 \right\} = \left(\mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^2}{\nu} \mathbf{I} \right)^{-1} \left(\mathbf{A}^\top \mathbf{y} + \frac{\sigma_y^2}{\nu} \bar{\mathbf{x}} \right). \quad (11)$$

Equation (11) can be computed using conjugate gradients (CG) or, if practical, the SVD $\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^\top$ via

$$\hat{\mathbf{x}} = \mathbf{V} \left(\mathbf{S}^\top \mathbf{S} + \frac{\sigma_y^2}{\nu} \mathbf{I} \right)^{-1} \left(\mathbf{S}^\top \mathbf{U}^\top \mathbf{y} + \frac{\sigma_y^2}{\nu} \mathbf{V}^\top \bar{\mathbf{x}} \right). \quad (12)$$

In any case, from (8) and (11), the error in $\hat{\mathbf{x}}$ can be written as

$$\hat{\mathbf{x}} - \mathbf{x}_0 = \left(\mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^2}{\nu} \mathbf{I} \right)^{-1} \left(\mathbf{A}^\top [\mathbf{x}_0 + \sigma_y \mathbf{w}] + \frac{\sigma_y^2}{\nu} [\mathbf{x}_0 - \sqrt{\nu} \mathbf{e}] \right) - \mathbf{x}_0 = \left(\mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^2}{\nu} \mathbf{I} \right)^{-1} \left(\sigma_y \mathbf{A}^\top \mathbf{w} - \frac{\sigma_y^2}{\sqrt{\nu}} \mathbf{e} \right), \quad (13)$$

and so the covariance of the error in $\hat{\mathbf{x}}$ can be written as

$$\text{Cov}\{\hat{\mathbf{x}} - \mathbf{x}_0\} = \left(\mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^2}{\nu} \mathbf{I} \right)^{-1} \left(\sigma_y^2 \mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^4}{\nu} \mathbf{I} \right) \left(\mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^2}{\nu} \mathbf{I} \right)^{-1} = \left(\frac{1}{\sigma_y^2} \mathbf{A}^\top \mathbf{A} + \frac{1}{\nu} \mathbf{I} \right)^{-1} \triangleq \mathbf{C}. \quad (14)$$

From (14) we see that the error in $\hat{\mathbf{x}}$ can be strongly colored. For example, in the case of inpainting, where \mathbf{A} is formed from rows of the identity matrix, the error variance in the masked pixels equals ν , while the error variance in the unmasked pixels equals $(1/\sigma_y^2 + 1/\nu)^{-1} \leq \sigma_y^2$. These two values may differ by many orders of magnitude. Since most denoisers are trained to remove white noise with a specified variance of σ^2 , direct denoising of $\hat{\mathbf{x}}$ performs poorly, as we show in Sec. 4.1.

To circumvent the issues that arise from colored denoiser-input error, we propose to add ‘‘complementary’’ colored Gaussian noise $\mathbf{c} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$ to $\hat{\mathbf{x}}$ so that the resulting $\mathbf{r} = \hat{\mathbf{x}} + \mathbf{c}$ has an error covariance of $\sigma^2 \mathbf{I}$, i.e., white error. This requires that

$$\mathbf{\Sigma} = \sigma^2 \mathbf{I} - \mathbf{C} = \sigma^2 \mathbf{I} - \left(\frac{1}{\sigma_y^2} \mathbf{V} \mathbf{S}^\top \mathbf{S} \mathbf{V}^\top + \frac{1}{\nu} \mathbf{I} \right)^{-1} = \mathbf{V} \text{Diag}(\boldsymbol{\lambda}) \mathbf{V}^\top \text{ for } \lambda_i = \sigma^2 - \frac{1}{s_i^2/\sigma_y^2 + 1/\nu} \quad (15)$$

for $s_i^2 \triangleq [\mathbf{S}^\top \mathbf{S}]_{i,i}$. By choosing $\sigma^2 \geq \nu$, we ensure that $\lambda_i \geq 0 \ \forall i$, needed for $\mathbf{\Sigma}$ to be a valid covariance matrix. In the case that the SVD is practical to implement, we can generate \mathbf{c} using

$$\mathbf{c} = \mathbf{V} \text{Diag}(\boldsymbol{\lambda})^{1/2} \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (16)$$

In the absence of an SVD, we propose to approximate $\mathbf{\Sigma}$ by

$$\hat{\mathbf{\Sigma}} \triangleq (\sigma^2 - \nu) \mathbf{I} + \xi \mathbf{A}^\top \mathbf{A} \quad (17)$$

with some $\xi \geq 0$. Note that $\hat{\mathbf{\Sigma}}$ agrees with $\mathbf{\Sigma}$ in the nullspace of \mathbf{A} (i.e., when $s_n = 0$) for any ξ . By choosing

$$\xi = \frac{1}{s_{\max}^2} \left(\nu - \frac{1}{s_{\max}^2/\sigma_y^2 + 1/\nu} \right), \quad (18)$$

$\hat{\mathbf{\Sigma}}$ will also agree with $\mathbf{\Sigma}$ in the strongest measured subspace (i.e., when $s_n = s_{\max}$). Without an SVD, s_{\max} can be computed using the power iteration (Parlett, 1998). Finally, $\mathbf{c} \sim \mathcal{N}(\mathbf{0}, \hat{\mathbf{\Sigma}})$ can be generated via

$$\mathbf{c} = [\sqrt{\sigma^2 - \nu} \mathbf{I} \quad \sqrt{\xi} \mathbf{A}^\top] \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \in \mathbb{R}^{d+m}. \quad (19)$$

Figure 6 shows a close agreement between the ideal and approximate renoised error spectra in practice. Next, we provide the main theoretical result on FIRE.

Theorem 1 *Suppose that, for any input $\mathbf{r} = \mathbf{x}_0 + \sigma \boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, the denoiser output $\mathbf{d}_\theta(\mathbf{r}, \sigma)$ has white Gaussian error with known variance $\nu < \sigma^2$ and independent of the noise \mathbf{w} in (6). Then if initialized using $\mathbf{r}_{\text{init}} = \mathbf{x}_0 + \sigma_{\text{init}} \boldsymbol{\epsilon}$ with arbitrarily large but finite σ_{init} and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, there exists a $\rho > 1$ under which the FIRE iteration S1)-S4) converges to the true \mathbf{x}_0 .*

Appendix E provides a proof. Note that a key assumption of Theorem 1 is that the denoiser output error is white and Gaussian. Because this may not hold in practice, we propose to replace S1) with a ‘‘stochastic denoising’’ step (21), in which AWGN is explicitly added to the denoiser output. As the AWGN variance increases, the denoiser output becomes closer to white and Gaussian but its signal-to-noise ratio (SNR) degrades. To balance these competing objectives, we propose to add AWGN with variance approximately

Algorithm 1 FIRE for the SLM: $\hat{\mathbf{x}} = \text{FIRE}_{\text{SLM}}(\mathbf{y}, \mathbf{A}, \sigma_y, \mathbf{r}_{\text{init}}, \sigma_{\text{init}}, N, \rho)$

Require: $\mathbf{y}, \mathbf{A}, s_{\text{max}}, \sigma_y, N, \rho > 1, \mathbf{r}_{\text{init}}, \sigma_{\text{init}}$. Also $\mathbf{A} = \mathbf{U} \text{Diag}(\mathbf{s}) \mathbf{V}^T$ if using SVD.

```

1:  $\mathbf{r} = \mathbf{r}_{\text{init}}$  and  $\sigma = \sigma_{\text{init}}$   $\triangleright$  Initialize
2: for  $n = 1, \dots, N$  do
3:    $\bar{\mathbf{x}} \leftarrow \mathbf{d}_{\theta}(\mathbf{r}, \sigma) + \sqrt{\hat{\nu}_{\phi}(\sigma)} \mathbf{v}, \mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$   $\triangleright$  Stochastic denoising
4:    $\nu \leftarrow (\|\mathbf{y} - \mathbf{A}\bar{\mathbf{x}}\|^2 - \sigma_y^2 m) / \|\mathbf{A}\|_F^2$   $\triangleright$  Error variance of  $\bar{\mathbf{x}}$ 
5:    $\hat{\mathbf{x}} \leftarrow \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|^2 / \sigma_y^2 + \|\mathbf{x} - \bar{\mathbf{x}}\|^2 / \nu$   $\triangleright$  Estimate  $\mathbf{x}_0 \sim \mathcal{N}(\bar{\mathbf{x}}, \nu \mathbf{I})$  from  $\mathbf{y} \sim \mathcal{N}(\mathbf{A}\mathbf{x}_0, \sigma_y^2 \mathbf{I})$ 
6:    $\sigma^2 \leftarrow \max\{\sigma^2 / \rho, \nu\}$   $\triangleright$  Decrease target variance
7:   if have SVD then
8:      $\lambda_i \leftarrow \sigma^2 - (s_i^2 / \sigma_y^2 + 1 / \nu)^{-1}, i = 1, \dots, d$ 
9:      $\mathbf{c} \leftarrow \mathbf{V} \text{Diag}(\boldsymbol{\lambda})^{1/2} \boldsymbol{\epsilon}, \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$   $\triangleright$  Colored Gaussian noise
10:  else
11:     $\xi \leftarrow (\nu - (s_{\text{max}}^2 / \sigma_y^2 + 1 / \nu)^{-1}) / s_{\text{max}}^2$ 
12:     $\mathbf{c} \leftarrow [\sqrt{\sigma^2 - \nu} \mathbf{I} \quad \sqrt{\xi} \mathbf{A}^T] \boldsymbol{\epsilon}, \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$   $\triangleright$  Colored Gaussian noise
13:   $\mathbf{r} \leftarrow \hat{\mathbf{x}} + \mathbf{c}$   $\triangleright$  Renoise so that  $\text{Cov}\{\mathbf{r} - \mathbf{x}_0\} = \sigma^2 \mathbf{I}$ 
14: return  $\hat{\mathbf{x}}$ 

```

equal to that of the raw-denoiser output error. We estimate the latter quantity from the denoiser input variance σ^2 by training a predictor of the form

$$\hat{\nu}_{\phi}(\sigma) \approx \mathbb{E}\{\|\mathbf{d}_{\theta}(\mathbf{x}_0 + \sigma \boldsymbol{\epsilon}, \sigma) - \mathbf{x}_0\|^2 / d\}, \quad (20)$$

where the expectation is over $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and validation images $\mathbf{x}_0 \sim p_0$. Recall that d is the dimension of \mathbf{x}_0 . In our experiments, $\hat{\nu}_{\phi}(\cdot)$ is implemented using a lookup table. The stochastic denoising step is then

$$\bar{\mathbf{x}} = \mathbf{d}_{\theta}(\mathbf{x}, \sigma) + \sqrt{\hat{\nu}_{\phi}(\sigma)} \mathbf{v}, \mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (21)$$

Algorithm 1 summarizes the FIRE algorithm for the SLM (6). In App. D, we describe a minor enhancement to Alg. 1 that speeds up the MMSE estimation step when CG is used.

3.2 DDfire for the GLM

We now propose to extend the SLM-FIRE from Sec. 3.1 to the generalized linear model (GLM)

$$\mathbf{y} \sim p(\mathbf{y} | \mathbf{z}_0) = \prod_{j=1}^m p_{y|z}(y_j | z_{0,j}) \quad \text{with} \quad \mathbf{z}_0 \triangleq \mathbf{A}\mathbf{x}_0 \quad (22)$$

where $p_{y|z}$ is some scalar “measurement channel.” Examples include $p_{y|z}(y|z) = \mathcal{N}(y; |z|, \sigma_y^2)$ for phase retrieval, $p_{y|z}(y|z) = z^y e^{-z} / y!$ for Poisson regression, and $p_{y|z}(y|z) = \int_{\tau_y}^{\tau_y+1} \mathcal{N}(\tau; z, \sigma_y^2) d\tau$ for dequantization.

Our extension is inspired by expectation propagation (EP) (Minka, 2001; Bishop, 2007) and its application to GLMs (Schniter et al., 2016; Meng et al., 2018). The idea is to iterate between i) constructing “pseudo-measurements” $\bar{\mathbf{y}} = \mathbf{A}\mathbf{x}_0 + \bar{\mathbf{w}}$ with $\bar{\mathbf{w}} \sim \mathcal{N}(\mathbf{0}, \bar{\sigma}_y^2 \mathbf{I})$ using $p_{y|z}$ and an SLM-FIRE-constructed belief that $\mathbf{z}_0 \sim \mathcal{N}(\bar{\mathbf{z}}_0, \bar{\nu}_z \mathbf{I})$, and then ii) running SLM-FIRE with those pseudo-measurements and updating its belief on \mathbf{z}_0 . Figure 2 shows a high-level summary. Details are given below.

To construct the belief on \mathbf{z}_0 , we use the SLM-FIRE denoiser output model $\mathbf{x}_0 \sim \mathcal{N}(\bar{\mathbf{x}}, \nu \mathbf{I})$ from (8). Because $\mathbf{z}_0 = \mathbf{A}\mathbf{x}_0$, we see that $\mathbf{z}_0 \sim \mathcal{N}(\bar{\mathbf{z}}, \nu \mathbf{A} \mathbf{A}^T)$, where $\bar{\mathbf{z}} \triangleq \mathbf{A} \bar{\mathbf{x}}$. For simplicity, however, we use the white-noise approximation $\mathbf{z}_0 \sim \mathcal{N}(\bar{\mathbf{z}}, \bar{\nu}_z \mathbf{I})$, where $\bar{\nu}_z \triangleq \nu \|\mathbf{A}\|_F^2 / m$. Using the scalar belief $z_{0,j} \sim \mathcal{N}(\bar{z}_j, \bar{\nu}_z)$ and the likelihood model $y_j \sim p_{y|z}(\cdot | z_{0,j})$, EP suggests to first compute the posterior mean $\mathbb{E}\{z_{0,j} | y_j; \bar{z}_j, \bar{\nu}_z\} \triangleq \hat{z}_j$ and variance $\frac{1}{m} \sum_{j=1}^m \text{var}\{z_{0,j} | y_j; \bar{z}_j, \bar{\nu}_z\} \triangleq \hat{\nu}_z$, and then pass the “extrinsic” versions of those quantities:

$$\bar{\sigma}_y^2 \triangleq [1 / \hat{\nu}_z - 1 / \bar{\nu}_z]^{-1}, \quad \bar{\mathbf{y}} \triangleq \bar{\sigma}_y^2 (\hat{\mathbf{z}} / \hat{\nu}_z - \bar{\mathbf{z}} / \bar{\nu}_z) \quad (23)$$




Figure 2: High-level overview of GLM-FIRE, which uses EP-style iterations between SLM-FIRE and an MMSE inference stage that involves the scalar measurement channel $p_{y|z}$.

back to SLM-FIRE, where they are used to construct the pseudo-measurement model

$$\bar{\mathbf{y}} = \mathbf{A}\mathbf{x}_0 + \bar{\sigma}_y \bar{\mathbf{w}}, \quad \bar{\mathbf{w}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (24)$$

In Fig. 9, we show that GLM-FIRE’s $\bar{\sigma}_y^2$ accurately tracks the true noise power in $\bar{\mathbf{y}}$, i.e., $\|\bar{\mathbf{y}} - \mathbf{A}\mathbf{x}_0\|^2/m$.

The GLM-FIRE algorithm is summarized as Alg. 3. When $p_{y|z}(y|z) = \mathcal{N}(y; z, \sigma_y^2)$, is it straightforward to show that $\bar{\mathbf{y}} = \mathbf{y}$ and $\bar{\sigma}_y^2 = \sigma_y^2$ for any $\bar{\mathbf{z}}$ and \bar{v}_z , in which case GLM-FIRE reduces to SLM-FIRE.

3.3 Putting FIRE into diffusion

Sections 3.1 and 3.2 detailed the FIRE algorithms for the SLM (6) and the GLM (22), respectively. In both cases, the FIRE algorithm approximates $\mathbb{E}\{\mathbf{x}_0|\mathbf{r}, \mathbf{y}\}$ given the measurements \mathbf{y} and the side-information $\mathbf{r} = \mathbf{x}_0 + \sigma\boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Thus, recalling the discussion in Sec. 2, FIRE can be used in the SMLD reverse process as an approximation of $\mathbb{E}\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\}$ by setting $\mathbf{r} = \mathbf{x}_t$ and $\sigma = \sigma_t$.

Instead of using SMLD for the diffusion reverse process, however, we use DDIM from Song et al. (2021a), which can be considered as a generalization of SMLD. In the sequel, we distinguish the DDIM quantities by writing them with subscript k . As detailed in App. C, DDIM is based on the model

$$\mathbf{x}_k = \mathbf{x}_0 + \sigma_k \boldsymbol{\epsilon}_k, \quad \boldsymbol{\epsilon}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (25)$$

for $k = 1, \dots, K$, where $\{\sigma_k^2\}_{k=1}^K$ is a specified sequence of variances. The DDIM reverse process iterates

$$\mathbf{x}_{k-1} = h_k \mathbf{x}_k + g_k \mathbb{E}\{\mathbf{x}_0|\mathbf{x}_k, \mathbf{y}\} + \varsigma_k \mathbf{n}_k \quad (26)$$

$$\varsigma_k = \eta_{\text{ddim}} \sqrt{\frac{\sigma_{k-1}^2(\sigma_k^2 - \sigma_{k-1}^2)}{\sigma_k^2}}, \quad h_k = \sqrt{\frac{\sigma_{k-1}^2 - \varsigma_k^2}{\sigma_k^2}}, \quad g_k = 1 - h_k \quad (27)$$

over $k = K, \dots, 2, 1$, starting from $\mathbf{x}_K \sim \mathcal{N}(\mathbf{0}, \sigma_K^2 \mathbf{I})$, using i.i.d $\{\mathbf{n}_k\}_{k=1}^K \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and some $\eta_{\text{ddim}} \geq 0$. When $\eta_{\text{ddim}} = 1$ and $K = T$, DDIM reduces to SMLD. But when $\eta_{\text{ddim}} = 0$, the DDIM reverse process (26) is deterministic and can be considered as a discretization of the probability-flow ODE Song et al. (2021a), which can outperform SMLD when the number of discretization steps K is small (Chen et al., 2023).

For a specified number K of DDIM steps (which we treat as a tuning parameter), we set the DDIM variances $\{\sigma_k^2\}_{k=1}^K$ as the geometric sequence

$$\sigma_k^2 = \sigma_{\min}^2 \left(\frac{\sigma_{\max}^2}{\sigma_{\min}^2} \right)^{\frac{k-1}{K-1}}, \quad k = 1, \dots, K \quad (28)$$

for some σ_{\min}^2 and σ_{\max}^2 that are typically chosen to match the minimum and maximum variances used to train the denoiser $\mathbf{d}_{\boldsymbol{\theta}}(\cdot, \cdot)$ or score approximation $\mathbf{s}_{\boldsymbol{\theta}}(\cdot, \cdot)$. So for example, if $\mathbf{s}_{\boldsymbol{\theta}}(\cdot, \cdot)$ was trained over the DDPM steps $t \in \{1, \dots, T\}$ for $T = 1000$, then we would set $\sigma_{\min}^2 = (1 - \bar{\alpha}_1)/\bar{\alpha}_1$ and $\sigma_{\max}^2 = (1 - \bar{\alpha}_{1000})/\bar{\alpha}_{1000}$ with $\bar{\alpha}_t$; see (37) for additional details.

Next we discuss how we set the FIRE iteration schedule $\{N_k\}_{k=1}^K$ and variance-decrease-factor $\rho > 1$. In doing so, we have two main goals:

- G1) Ensure that, at every DDIM step k , the denoiser’s output-error variance is at most ν_{thresh} at the final FIRE iteration, where ν_{thresh} is some value to be determined.




Figure 3: For an FFHQ denoiser: the geometric DDIM variances $\{\sigma_k^2\}_{k=1}^K$ versus DDIM step k for $K = 10$, the σ_{thresh}^2 corresponding to a $\delta = 0.4$ fraction of single-FIRE-iteration DDIM steps, and the denoiser input variance σ^2 at each FIRE iteration of each DDIM step, for $N_{\text{tot}} = 25$ total NFEs.

G2) Meet a fixed budget of $N_{\text{tot}} \triangleq \sum_{k=1}^K N_k$ total NFEs.

Note that, because the denoiser’s output-error variance increases monotonically with its input-error variance, we can rephrase G1) as

G1*) Ensure that, at every DDIM step k , the denoiser’s input-error variance is at most σ_{thresh}^2 at the final FIRE iteration, where σ_{thresh}^2 is some value to be determined.

Although σ_{thresh}^2 could be tuned directly, it’s not the most convenient option because a good search range can be difficult to construct. Instead, we tune the fraction $\delta \in [0, 1)$ of DDIM steps k that use a single FIRE iteration (i.e., that use $N_k = 1$) and we set σ_{thresh}^2 at the DDIM variance σ_k^2 of the first reverse-process step k that uses a single FIRE iteration, i.e., $1 + \lfloor (K - 1)\delta \rfloor \triangleq k_{\text{thresh}}$. (Note that $k_{\text{thresh}} = 1$ when $\delta = 0$ and $k_{\text{thresh}} = K - 1$ for $\delta \approx 1$.) All subsequent¹ DDIM steps $k < k_{\text{thresh}}$ will then automatically satisfy G1*) because σ_k^2 decreases with k .

To ensure that the earlier DDIM steps $k > k_{\text{thresh}}$ also satisfy G1*), we need that $\sigma_k^2 / \rho^{N_k - 1} \leq \sigma_{\text{thresh}}^2$, since σ_k^2 is the denoiser input-error variance at the first FIRE iteration and $\sigma_k^2 / \rho^{N_k - 1}$ is the denoiser input-error variance at the last FIRE iteration. For a fixed $\rho > 1$, we can rewrite this inequality as

$$N_k \geq \frac{\ln \sigma_k^2 - \ln \sigma_{\text{thresh}}^2}{\ln \rho} + 1 \triangleq \underline{N}_k. \quad (29)$$

Because N_k is a positive integer, it suffices to choose

$$N_k = \lceil \max\{1, \underline{N}_k\} \rceil \quad \forall k. \quad (30)$$

Finally, ρ is chosen as the smallest value that meets the NFE budget G2) under (30). We find this value using bisection search. For a given k_{thresh} , a lower bound on the total NFEs is $k_{\text{thresh}} \cdot 1 + (K - k_{\text{thresh}}) \cdot 2$. The definition of k_{thresh} then implies that $N_{\text{tot}} \geq K(2 - \delta) + \delta - 1$ and thus $K \leq (N_{\text{tot}} + 1 - \delta) / (2 - \delta) \triangleq K_{\text{min}}$.

In summary, for a budget of N_{tot} total NFEs, we treat the number of DDIM steps $K \in \{1, \dots, K_{\text{min}}\}$ and the fraction of single-FIRE-iteration steps $\delta \in [0, 1)$ as tuning parameters and, from them, compute $\{\sigma_k^2\}_{k=1}^K$, $\{N_k\}_{k=1}^K$, and ρ . Figure 3 shows an example. The pair (K, δ) can be tuned using cross-validation. Algorithm 2 details DDIM with the FIRE approximation of $E\{\mathbf{x}_0 | \mathbf{x}_k, \mathbf{y}\}$, which we refer to as “DDfire.”

¹Recall that the reverse process counts backwards, i.e., $k = K, K - 1, \dots, 2, 1$.

Algorithm 2 DDfire**Require:** $\mathbf{y}, \mathbf{A}, \sigma_y, \rho, \{\sigma_k\}_{k=1}^K, \{N_k\}_{k=1}^K, \eta_{\text{ddim}} \geq 0$ 1: $\mathbf{x}_K \sim \mathcal{N}(\mathbf{0}, \sigma_K^2 \mathbf{I})$ 2: **for** $k = K, K-1, \dots, 1$ **do**3: $\hat{\mathbf{x}}_{0|k} = \text{FIRE}(\mathbf{y}, \mathbf{A}, *, \mathbf{x}_k, \sigma_k, N_k, \rho)$ $\triangleright \text{FIRE via Alg. 1 or Alg. 3}$ 4: $\varsigma_k = \eta_{\text{ddim}} \sqrt{\frac{\sigma_{k-1}^2 (\sigma_k^2 - \sigma_{k-1}^2)}{\sigma_k^2}}$ 5: $\mathbf{x}_{k-1} = \sqrt{\frac{\sigma_{k-1}^2 - \varsigma_k^2}{\sigma_k^2}} \mathbf{x}_k + \left(1 - \sqrt{\frac{\sigma_{k-1}^2 - \varsigma_k^2}{\sigma_k^2}}\right) \hat{\mathbf{x}}_{0|k} + \varsigma_k \mathbf{n}_k, \quad \mathbf{n}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ $\triangleright \text{DDIM update}$ 6: **return** $\hat{\mathbf{x}}_{0|k}$ **3.4 Relation to other diffusion methods**

When estimating \mathbf{x}_0 from the measurements \mathbf{y} of (6), techniques like DDS (Chung et al., 2024) and DiffPIR (Zhu et al., 2023) also use a Gaussian prior approximation of the form $\mathbf{x}_0 \sim \mathcal{N}(\bar{\mathbf{x}}, \nu \mathbf{I})$. But while DDS treats ν as a tuning parameter and DiffPIR sets $\nu = \lambda \sigma_k^2$ for tunable λ , the proposed DDfire explicitly estimates ν using (10). Also, DDfire runs multiple inner iterations per diffusion step and employs colored renoising.

SLM-DDfire’s use of multiple denoising/renoising steps per DDIM iteration k bears some resemblance to the use of “time travel” in RePaint (Lugmayr et al., 2022). At step k of RePaint’s reverse process, the DDPM forward process is invoked to transform \mathbf{x}_k back to \mathbf{x}_{k+j} for $j \sim 10$ (which involves adding noise), after which j reverse steps are performed to obtain a new \mathbf{x}_k . This time traveling is repeated $r \sim 10$ times at each k , which has the effect of increasing the total NFEs by the factor r . While RePaint handles only inpainting problems, DDNM⁺ (Wang et al., 2023) extends its approach to general AWGN-corrupted linear inverse problems by using an SVD in a manner that closely resembles DDRM (Kawar et al., 2022a). In contrast, SLM-DDfire does not require an SVD and is designed to operate with many fewer NFEs.

Although we are unaware of prior work combining EP with diffusion, there is a paper by Meng & Kabashima (2024) that applies EP to the annealed Langevin dynamics approach from Song & Ermon (2020), although for the specific task of dequantization. Like with DPS and IIGDM, Bayes rule is used to write $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t | \mathbf{y}) = \nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t) + \nabla_{\mathbf{x}} \ln p_t(\mathbf{y} | \mathbf{x}_t)$ and the first term is approximated by $\mathbf{s}_{\theta}(\mathbf{x}_t, t)$ (recall Sec. 2). The second term is approximated with multiple EP iterations, each of which solves a linear system using an SVD. DDfire differs in using colored renoising, an SVD-free option, and DDIM diffusion, which allows far fewer NFEs.

4 Numerical experiments

We use 256×256 FFHQ (Karras et al., 2019) and ImageNet (Deng et al., 2009) datasets with pretrained diffusion models from Chung et al. (2023a) and Dhariwal & Nichol (2021), respectively. As linear inverse problems, we consider box inpainting with a 128×128 mask, Gaussian deblurring using a 61×61 blur kernel with 3-pixel standard deviation, motion deblurring using a 61×61 blur kernel generated using Borodenko (2020) with intensity 0.5, and $4\times$ bicubic super-resolution. We compare to DDRM (Kawar et al., 2022a), DiffPIR (Zhu et al., 2023), IIGDM (Song et al., 2023), DPS (Chung et al., 2023a), RED-diff (Mardani et al., 2024), and DAPS (Zhang et al., 2025).

We also consider phase retrieval with the shot-noise corruption mechanism from Metzler et al. (2018) for both oversampled Fourier (OSF) and coded diffraction pattern (CDP) (Candès et al., 2015) \mathbf{A} at $4\times$ oversampling with $\alpha_{\text{shot}} = 8$ and 45, respectively. Here, α_{shot} is the shot-noise strength, as detailed in App. G. We compare to prDeep (Metzler et al., 2018), DOLPH (Shoushtari et al., 2023), DPS, RED-diff, DAPS, and the classical hybrid input-output (HIO) algorithm (Fienup, 1982), using $p_{\mathbf{y}|\mathbf{z}}(\mathbf{y}|\mathbf{z}) = \mathcal{N}(\mathbf{y}; |\mathbf{z}|, \sigma_y^2)$ for all algorithms that accept a likelihood function.

Unless specified otherwise, DDfire was configured as follows. For the linear inverse problems, CG is used (no SVD), 1000 NFEs are used without stochastic denoising, and the (K, δ) hyperparameters are tuned

Table 1: DDfire ablation results for noisy FFHQ box inpainting with $\sigma_y = 0.05$ at 1000 NFEs.

Method	PSNR \uparrow	LPIPS \downarrow	Runtime
DDfire	24.31	0.1127	34.37s
DDfire w/o renoising	18.48	0.2349	34.37s
DDfire w/o colored renoising	23.64	0.1553	34.37s
DDfire w/ stochastic denoising	24.30	0.1143	34.37s
DDfire w/o estimating ν	23.02	0.1755	34.37s
DDfire w/o CG early stopping	24.31	0.1127	52.12s
DDfire w/ SVD	24.31	0.1124	30.97s

Table 2: Noisy FFHQ results with measurement noise standard deviation $\sigma_y = 0.05$.

Model	Inpaint (box)			Deblur (Gaussian)			Deblur (Motion)			4 \times Super-resolution		
	PSNR \uparrow	LPIPS \downarrow	FID \downarrow	PSNR \uparrow	LPIPS \downarrow	FID \downarrow	PSNR \uparrow	LPIPS \downarrow	FID \downarrow	PSNR \uparrow	LPIPS \downarrow	FID \downarrow
DDRM	21.71	0.1551	40.61	25.35	0.2223	51.70	-	-	-	27.32	0.1864	45.82
DiffPIR	22.43	0.1883	31.98	24.56	0.2394	34.82	26.91	0.1952	26.67	24.89	0.2486	32.33
IIGDM	21.41	0.2009	44.41	23.66	0.2525	45.34	25.14	0.2082	41.95	24.40	0.2520	51.41
DPS	22.54	0.1368	35.69	25.70	0.1774	25.18	26.74	0.1655	27.17	26.30	0.1850	27.38
RED-diff	23.58	0.1883	48.86	26.99	0.2081	38.82	16.47	0.5074	128.68	25.61	0.3569	70.86
DAPS	23.61	0.1415	31.51	26.97	0.1827	31.10	27.13	0.1718	30.74	26.91	0.1885	30.83
DDfire	24.75	0.1101	25.26	27.10	0.1533	24.97	28.14	0.1374	26.12	27.13	0.1650	25.73

to minimize LPIPS (Zhang et al., 2018) on a 100-sample validation set (see Table 5). For phase retrieval, 800 NFEs are used (to fairly compare to prDeep) with stochastic denoising and (K, δ) were the hand-tuned values in Table 6. Neither CG nor SVD are needed since (11) can be solved in analytically. Appendix G contains additional details on the implementation of DDfire and the competing methods.

4.1 Ablation study

We first perform an ablation study on the SLM-DDfire design choices in Sec. 3 using noisy FFHQ box inpainting and a 100-image validation set. The results are summarized in Table 1. We first see that both PSNR and LPIPS suffer significantly when FIRE is run without renoising (i.e., $\mathbf{c} = \mathbf{0}$ in line 13 in Alg. 1). Similarly, renoising using white noise (i.e., $\mathbf{c} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ in line 13 of Alg. 1) gives noticeably worse PSNR and LPIPS than the proposed colored noise. Using stochastic denoising gives nearly identical performance to plain denoising (i.e., $\hat{\nu}_\phi(\sigma) = 0$ in line 3 of Alg. 1), and so we use plain denoising by default with linear inverse problems. A more significant degradation results when the denoiser output-error variance ν is not adapted to $\bar{\mathbf{x}}$ in line 4 of Alg. 1 but set at the data-average value $\hat{\nu}_\phi(\sigma)$. On the other hand, when CG doesn’t use early stopping (as described in App. D) the runtime increases without improving PSNR or LPIPS. Thus, we use early stopping by default. Finally, using an SVD instead of CG, which also avoids the noise approximation in (17), gives essentially identical PSNR and LPIPS but with a slightly faster runtime. Figure 4 shows another LPIPS/runtime comparison of the SVD and CG versions of DDfire.

4.2 PSNR, LPIPS, and FID results

For noisy linear inverse problems, Tables 2–3 show PSNR, LPIPS, and FID (Heusel et al., 2017) on a 1000-sample test set for FFHQ and ImageNet data, respectively. DDRM was not applied to motion deblurring due to the lack of an SVD. Tables 2–3 show that DDfire wins in most cases and otherwise performs well.

Fig. 5 shows image examples for inpainting, motion deblurring, Gaussian deblurring, and 4 \times super-resolution on ImageNet. The zoomed regions show that DDfire did a better job recovering fine details. Additional examples can be found in Fig. 11.

For OSF and CDP phase retrieval, Table 4 shows PSNR, LPIPS, and FID on a 1000-sample test set for FFHQ. For both CDP and OSF, the table shows that DDfire outperformed the competitors in all metrics.

Table 3: Noisy ImageNet results with measurement noise standard deviation $\sigma_y = 0.05$.

Model	Inpaint (box)			Deblur (Gaussian)			Deblur (Motion)			4× Super-resolution		
	PSNR↑	LPIPS↓	FID↓	PSNR↑	LPIPS↓	FID↓	PSNR↑	LPIPS↓	FID↓	PSNR↑	LPIPS↓	FID↓
DDRM	18.24	0.2423	67.47	22.56	0.3454	68.78	-	-	-	24.49	0.2777	64.68
DiffPIR	18.03	0.2860	65.55	21.31	0.3683	56.35	24.36	0.2888	54.11	23.31	0.3383	63.48
PIGDM	17.69	0.3303	86.36	20.87	0.4191	75.43	22.15	0.3591	70.91	21.25	0.4149	78.57
DPS	18.23	0.2314	59.10	21.30	0.3393	50.46	21.77	0.3307	80.27	23.38	0.2904	49.86
RED-diff	18.95	0.2909	108.88	23.45	0.3190	65.65	15.21	0.5647	198.74	22.99	0.3858	83.06
DAPS	19.99	0.2199	61.53	23.91	0.2863	56.87	24.58	0.2722	54.83	24.04	0.2729	55.54
DDfire	20.39	0.1915	55.54	23.71	0.2353	50.05	24.59	0.2314	49.25	23.58	0.2629	49.67

Table 4: Noisy FFHQ phase retrieval results

Model	OSF			CDP		
	PSNR↑	LPIPS↓	FID↓	PSNR↑	LPIPS↓	FID↓
HIO	23.66	0.4706	130.58	17.59	0.5430	84.87
DOLPH	14.73	0.7220	389.88	25.76	0.1686	32.93
DPS	23.63	0.2908	53.91	29.19	0.1394	27.87
RED-diff	25.47	0.2828	65.74	28.75	0.1734	28.87
DAPS	24.10	0.2891	57.73	28.26	0.1927	34.97
prDeep	30.90	0.1132	31.51	19.24	0.4183	59.44
DDfire	33.56	0.0691	28.94	30.01	0.1232	23.49




Figure 4: LPIPS vs. single image sampling time for noisy Gaussian deblurring on an A100 GPU. The evaluation used 1000 ImageNet images. Solid line: DDfire with CG for various numbers of NFEs. Dashed line: DDfire with SVD.

For OSF phase retrieval, DDfire outperformed the best competitor (prDeep) by 2.6 dB and the second best (RED-diff) by 8.1 dB in PSNR. Example reconstructions can be found in Fig. 10.

4.3 Runtime results

Figure 4 shows LPIPS vs. average runtime (in seconds on an A100 GPU) to generate a single image for noisy Gaussian deblurring on the 1000-sample ImageNet test set. The figure shows that DDfire gives a better performance/complexity tradeoff than the competitors. It also shows that DDfire is approximately 1.5 times faster than DPS when both are run at 1000 NFEs, due to DPS’s use of backpropagation.

Figure 7 shows LPIPS vs. runtime for noisy OSF phase retrieval on the 1000-sample FFHQ test set, showing that DDfire gives a significantly better performance/complexity tradeoff than all diffusion-based competitors.




Figure 5: Example recoveries from noisy linear inverse problems with ImageNet images.

5 Conclusion

To solve linear inverse problems, we proposed the Fast Iterative Renoising (FIRE) algorithm, which can be interpreted as the HQS plug-and-play algorithm with a colored renoising step that aims to whiten the denoiser input error. We then extended the linear FIRE algorithm to the generalized-linear case using expectation propagation (EP). Since these FIRE algorithms approximate the measurement-conditional denoiser $E\{\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}\}$, or equivalently the measurement-conditional score $\nabla_{\mathbf{x}} \ln p_t(\mathbf{x}_t|\mathbf{y})$, they can be readily combined with DDIM for diffusion posterior sampling, giving the “DDfire” algorithm. Experiments on box inpainting, Gaussian and motion deblurring, and $4\times$ super-resolution with FFHQ and ImageNet images show DDfire outperforming DDRM, IIGDM, DiffPIR, DPS, RED-diff, and DAPS in PSNR, LPIPS, and FID metrics in nearly all cases. Experiments on noisy FFHQ phase retrieval (both OSF and CDP versions) show DDfire outperforming HIO, prDeep, DOLPH, DPS, RED-diff, and DAPS in all cases. Finally, DDfire offers fast inference, with better LPIPS-versus-runtime curves than the competitors.

References

- R. Ahmad, C. A. Bouman, G. T. Buzzard, S. Chan, S. Liu, E. T. Reehorst, and P. Schniter. Plug and play methods for magnetic resonance imaging. *IEEE Signal Process. Mag.*, 37(1):105–116, March 2020.
- Tamir Bendory, Robert Beinert, and Yonina C Eldar. Fourier phase retrieval: Uniqueness and algorithms. In *MATHEON Conf. Compressed Sensing and its Applications*, pp. 55–91, 2015.
- C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, New York, 2007.
- Levi Borodenko. motionblur. Downloaded from <https://github.com/LeviBorodenko/motionblur>, 2020.
- E. J. Candès, X. Li, and M. Soltanolkotabi. Phase retrieval from coded diffraction patterns. *Appl. Comput. Harmonic Anal.*, 39(2):277–299, 2015.
- Sitan Chen, Sinho Chewi, Holden Lee, Yuanzhi Li, Jianfeng Lu, and Adil Salim. The probability flow ODE is provably fast. In *Proc. Neural Info. Process. Syst. Conf.*, volume 36, pp. 68552–68575, 2023.
- Hyungjin Chung, Jeongsol Kim, Michael T McCann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *Proc. Intl. Conf. Learn. Rep.*, 2023a.
- Hyungjin Chung, Jeongsol Kim, Michael T McCann, Marc L Klasky, and Jong Chul Ye. diffusion-posterior-sampling. <https://github.com/DPS2022/diffusion-posterior-sampling>, March 2023b.
- Hyungjin Chung, Suhyeon Lee, and Jong Chul Ye. Decomposed diffusion sampler for accelerating large-scale inverse problems. In *Proc. Intl. Conf. Learn. Rep.*, 2024.
- Florentin Coeurdoux, Nicolas Dobigeon, and Pierre Chainais. Plug-and-play split Gibbs sampler: Embedding deep generative priors in Bayesian inference. *IEEE Trans. Image Process.*, 33:3496–3507, 2024.
- Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar, Alexandros G Dimakis, and Mauricio Delbracio. A survey on diffusion models for inverse problems. *arXiv:2410.00083*, 2024.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. IEEE Conf. Comp. Vision Pattern Recog.*, pp. 248–255, 2009.
- P. Dhariwal and A. Nichol. Diffusion models beat GANs on image synthesis. In *Proc. Neural Info. Process. Syst. Conf.*, volume 34, pp. 8780–8794, 2021.
- B. Efron. Tweedie’s formula and selection bias. *J. Am. Statist. Assoc.*, 106(496):1602–1614, 2011.
- J. R. Fienup. Phase retrieval algorithms: A comparison. *Appl. Optics*, 21(15):2758–2769, August 1982.
- M. A. T. Figueiredo and J. M Bioucas-Dias. Restoration of Poissonian images using alternating direction optimization. *IEEE Trans. Image Process.*, 19(12):3133–3145, 2010.
- Doeke Hekstra, Ian Hunt-Isaak, Jack Greisman, and John Russell. phase-retrieval. <https://github.com/Hekstra-Lab/phase-retrieval>, 2018.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Proc. Neural Info. Process. Syst. Conf.*, volume 30, 2017.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proc. Neural Info. Process. Syst. Conf.*, volume 33, pp. 6840–6851, 2020.
- Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *J. Mach. Learn. Res.*, 6:695–709, 2005.
- Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proc. IEEE Conf. Comp. Vision Pattern Recog.*, pp. 4396–4405, 2019.

- Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *Proc. Neural Info. Process. Syst. Conf.*, 2022a.
- Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. Downloaded from <https://github.com/bahjat-kawar/ddrm>, May 2022b.
- David G Luenberger and Yinyu Ye. *Linear and Nonlinear Programming*. Springer, 2016.
- Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proc. IEEE Conf. Comp. Vision Pattern Recog.*, pp. 11461–11471, 2022.
- Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving inverse problems with diffusion models. In *Proc. Intl. Conf. Learn. Rep.*, 2024.
- Xiangming Meng and Yoshiyuki Kabashima. QCS-GGM+: Improved quantized compressed sensing with score-based generative models. In *Proc. AAAI Conf. Artificial Intell.*, volume 38, pp. 14341–14349, 2024.
- Xiangming Meng, Sheng Wu, and Jiang Zhu. A unified Bayesian inference framework for generalized linear models. *IEEE Signal Process. Lett.*, 25(3):398–402, 2018.
- C. A. Metzler. prdeep. <https://github.com/ricedsp/prDeep/tree/master>, 2018.
- C. A. Metzler, P. Schniter, A. Veeraraghavan, and R. G. Baraniuk. prDeep: Robust phase retrieval with flexible deep neural networks. In *Proc. Intl. Conf. Mach. Learn.*, pp. 3501–3510, 2018.
- T. Minka. *A Family of Approximate Algorithms for Bayesian Inference*. PhD thesis, Dept. Comp. Sci. Eng., MIT, Cambridge, MA, USA, 2001.
- NVlabs. RED-diff. Downloaded from <https://github.com/NVlabs/RED-diff>, 2023.
- Beresford N Parlett. *The symmetric eigenvalue problem*. SIAM, 1998.
- H. V. Poor. *An Introduction to Signal Detection and Estimation*. Springer, New York, NY, USA, 2nd edition, 1994.
- E. T. Reehorst and P. Schniter. Regularization by denoising: Clarifications and new interpretations. *IEEE Trans. Comput. Imag.*, 5(1):52–67, March 2019.
- Yaniv Romano, Michael Elad, and Peyman Milanfar. The little engine that could: Regularization by denoising (RED). *SIAM J. Imag. Sci.*, 10(4):1804–1844, 2017.
- Philip Schniter, Sundeep Rangan, and Alyson K Fletcher. Vector approximate message passing for the generalized linear model. In *Proc. Asilomar Conf. Signals Syst. Comput.*, pp. 1525–1529, 2016.
- Yoav Shechtman, Yonina C Eldar, Oren Cohen, Henry Nicholas Chapman, Jianwei Miao, and Mordechai Segev. Phase retrieval with application to optical imaging: A contemporary overview. *IEEE Signal Process. Mag.*, 32(3):87–109, 2015.
- Shirin Shoushtari, Jiaming Liu, and Ulugbek S Kamilov. Diffusion models for phase retrieval in computational imaging. In *Proc. Asilomar Conf. Signals Syst. Comput.*, pp. 779–783, 2023.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proc. Intl. Conf. Mach. Learn.*, pp. 2256–2265, 2015.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Proc. Intl. Conf. Learn. Rep.*, 2021a.
- Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *Proc. Intl. Conf. Learn. Rep.*, 2023.

- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Proc. Neural Info. Process. Syst. Conf.*, 2019.
- Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In *Proc. Neural Info. Process. Syst. Conf.*, 2020.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Proc. Intl. Conf. Learn. Rep.*, 2021b.
- Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. In *Proc. Intl. Conf. Learn. Rep.*, 2023.
- Zihui Wu, Yu Sun, Yifan Chen, Bingliang Zhang, Yisong Yue, and Katherine Bouman. Principled probabilistic imaging using diffusion models as plug-and-play priors. In *Proc. Neural Info. Process. Syst. Conf.*, 2024.
- Xingyu Xu and Yuejie Chi. Provably robust score-based diffusion posterior sampling for plug-and-play image reconstruction. In *Proc. Neural Info. Process. Syst. Conf.*, 2024.
- Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. Improving diffusion inverse problem solving with decoupled noise annealing. Downloaded from <https://github.com/zhangbingliang2019/DAPS>, 2024.
- Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. Improving diffusion inverse problem solving with decoupled noise annealing. In *Proc. IEEE Conf. Comp. Vision Pattern Recog.*, 2025.
- Kai Zhang, Yawei Li, Wangmeng Zuo, Lei Zhang, Luc Van Gool, and Radu Timofte. Plug-and-play image restoration with deep denoiser prior. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(10):6360–6376, 2021.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. IEEE Conf. Comp. Vision Pattern Recog.*, pp. 586–595, 2018.
- Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc Van Gool. Denoising diffusion models for plug-and-play image restoration. In *Proc. IEEE Conf. Comp. Vision Pattern Recog.*, pp. 1219–1229, 2023.
- Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc Van Gool. Diffpir. Downloaded from <https://github.com/yuanzhi-zhu/DiffPIR>, July 2024.
- A. Zymnis, S. Boyd, and E Candès. Compressed sensing with quantized measurements. *IEEE Signal Process. Lett.*, 17(2):149–152, 2010.

A VP formulation

In the main text, we describe DDfire for the VE SDE formulation from Song et al. (2021b) and the corresponding SMLD discretization from Song & Ermon (2019). Here, we describe it for the VP SDE from Song et al. (2021b) and the corresponding DDPM discretization from Ho et al. (2020).

From Song et al. (2021b), the general SDE forward process can be written as

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, t) dt + g(t) d\mathbf{w} \quad (31)$$

for some choices of $\mathbf{f}(\cdot, \cdot)$ and $g(\cdot)$, where $d\mathbf{w}$ is the standard Wiener process (i.e., Brownian motion). The reverse process can then be described by

$$d\mathbf{x} = (\mathbf{f}(\mathbf{x}, t) - g^2(t)\nabla_{\mathbf{x}} \ln p_t(\mathbf{x})) dt + g(t) d\bar{\mathbf{w}}, \quad (32)$$

where $p_t(\cdot)$ is the distribution of \mathbf{x} at time t and $d\bar{\mathbf{w}}$ is the reverse Wiener process. In the VE-SDE, $\mathbf{f}(\mathbf{x}, t) = \mathbf{0}$ and $g(t) = \sqrt{d[\sigma^2(t)]/dt}$ for some variance schedule $\sigma^2(t)$, but in the VP-SDE, $\mathbf{f}(\mathbf{x}, t) = -\frac{1}{2}\beta(t)\mathbf{x}$ and $g(t) = \sqrt{\beta(t)}$ for some variance schedule $\beta(t)$. When discretized to $t \in \{0, 1, \dots, T\}$, the VP forward process becomes

$$\tilde{\mathbf{x}}_t = \sqrt{1 - \beta_t} \tilde{\mathbf{x}}_{t-1} + \sqrt{\beta_t} \tilde{\mathbf{w}}_{t-1} \quad (33)$$

with i.i.d. $\{\tilde{\mathbf{w}}_t\} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, so that

$$\tilde{\mathbf{x}}_t = \sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} \tilde{\boldsymbol{\epsilon}}_t \quad (34)$$

with $\alpha_t \triangleq 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, and $\tilde{\boldsymbol{\epsilon}}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Throughout, we write the VP quantities with tildes to distinguish them from the VE quantities. The DDPM reverse process then takes the form

$$\tilde{\mathbf{x}}_{t-1} = \frac{1}{\sqrt{\alpha_t}} (\tilde{\mathbf{x}}_t + \beta_t \nabla_{\mathbf{x}_t} \ln p(\tilde{\mathbf{x}}_t)) + \Sigma_t \tilde{\mathbf{n}}_t \quad \text{with} \quad \Sigma_t^2 \triangleq \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t \quad (35)$$

and is typically initialized at $\tilde{\mathbf{x}}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. By rewriting (34) as

$$\frac{1}{\sqrt{\alpha_t}} \tilde{\mathbf{x}}_t = \mathbf{x}_0 + \sqrt{\frac{1 - \bar{\alpha}_t}{\alpha_t}} \tilde{\boldsymbol{\epsilon}}_t \quad (36)$$

and comparing it to (3), we recognize the VP/VE relationships

$$\frac{1}{\sqrt{\alpha_t}} \tilde{\mathbf{x}}_t = \mathbf{x}_t \quad \text{and} \quad \frac{1 - \bar{\alpha}_t}{\alpha_t} = \sigma_t^2 \Leftrightarrow \bar{\alpha}_t = \frac{1}{1 + \sigma_t^2}. \quad (37)$$

Furthermore, assuming that $\bar{\alpha}_T \ll 1$, the VP initialization $\tilde{\mathbf{x}}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is well approximated by the VE initialization $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \sigma_T^2 \mathbf{I})$.

B DDIM details for VP

The DDIM reverse process from Song et al. (2021a) provides an alternative to the DDPM reverse process that offers a flexible level of stochasticity. When describing VP DDIM, we will write the quantities as $\tilde{\mathbf{x}}_k, \tilde{\boldsymbol{\epsilon}}_k, \tilde{\mathbf{n}}_k, \bar{\alpha}_k$ to distinguish them from the corresponding VP DDPM quantities $\tilde{\mathbf{x}}_t, \tilde{\boldsymbol{\epsilon}}_t, \tilde{\mathbf{n}}_t, \bar{\alpha}_t$, and we will write the total number of steps as K . Like (34), DDIM is built around the model

$$\tilde{\mathbf{x}}_k = \sqrt{\alpha_k} \mathbf{x}_0 + \sqrt{1 - \alpha_k} \tilde{\boldsymbol{\epsilon}}_k, \quad \tilde{\boldsymbol{\epsilon}}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (38)$$

Adapting the first two equations from Song et al. (2021a, App.D.3) to our notation, we have

$$\tilde{\mathbf{x}}_{k-1} = \sqrt{\alpha_{k-1}} \left(\frac{\tilde{\mathbf{x}}_k - \sqrt{1 - \bar{\alpha}_k} \mathbb{E}\{\tilde{\boldsymbol{\epsilon}}_k | \tilde{\mathbf{x}}_k, \mathbf{y}\}}{\sqrt{\alpha_k}} \right) + \sqrt{1 - \alpha_{k-1} - \zeta_k^2} \mathbb{E}\{\tilde{\boldsymbol{\epsilon}}_k | \tilde{\mathbf{x}}_k, \mathbf{y}\} + \zeta_k \tilde{\mathbf{n}}_k \quad (39)$$

$$\zeta_k \triangleq \eta_{\text{ddim}} \sqrt{\frac{1 - \bar{\alpha}_{k-1}}{1 - \bar{\alpha}_k}} \sqrt{1 - \frac{\bar{\alpha}_k}{\alpha_{k-1}}} \quad (40)$$

with $\tilde{\mathbf{n}}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and tunable $\eta_{\text{ddim}} \geq 0$. When $\eta_{\text{ddim}} = 1$ and $K = T$, DDIM reduces to DDPM. But when $\eta_{\text{ddim}} = 0$, the reverse process is deterministic. In fact, it can be considered a discretization of the probability flow ODE (Song et al., 2021a), which often works much better than the SDE when the number of discretization steps K is small. We now write (39) in a simpler form. Applying $\mathbb{E}\{\cdot|\tilde{\mathbf{x}}_k, \mathbf{y}\}$ to both sides of (38) gives

$$\tilde{\mathbf{x}}_k = \sqrt{\bar{\alpha}_k} \mathbb{E}\{\mathbf{x}_0|\tilde{\mathbf{x}}_k, \mathbf{y}\} + \sqrt{1 - \bar{\alpha}_k} \mathbb{E}\{\tilde{\epsilon}_k|\tilde{\mathbf{x}}_k, \mathbf{y}\} \Leftrightarrow \mathbb{E}\{\tilde{\epsilon}_k|\tilde{\mathbf{x}}_k, \mathbf{y}\} = \frac{\tilde{\mathbf{x}}_k - \sqrt{\bar{\alpha}_k} \mathbb{E}\{\mathbf{x}_0|\tilde{\mathbf{x}}_k, \mathbf{y}\}}{\sqrt{1 - \bar{\alpha}_k}}, \quad (41)$$

and plugging (41) into (39) gives

$$\tilde{\mathbf{x}}_{k-1} = \sqrt{\bar{\alpha}_{k-1}} \mathbb{E}\{\mathbf{x}_0|\tilde{\mathbf{x}}_k, \mathbf{y}\} + \tilde{\varsigma}_k \tilde{\mathbf{n}}_k + \sqrt{1 - \bar{\alpha}_{k-1} - \tilde{\varsigma}_k^2} \left(\frac{\tilde{\mathbf{x}}_k - \sqrt{\bar{\alpha}_k} \mathbb{E}\{\mathbf{x}_0|\tilde{\mathbf{x}}_k, \mathbf{y}\}}{\sqrt{1 - \bar{\alpha}_k}} \right) \quad (42)$$

$$= \sqrt{\bar{\alpha}_{k-1}} \mathbb{E}\{\mathbf{x}_0|\tilde{\mathbf{x}}_k, \mathbf{y}\} + \tilde{\varsigma}_k \tilde{\mathbf{n}}_k + \tilde{h}_k (\tilde{\mathbf{x}}_k - \sqrt{\bar{\alpha}_k} \mathbb{E}\{\mathbf{x}_0|\tilde{\mathbf{x}}_k, \mathbf{y}\}) \quad (43)$$

$$= \tilde{h}_k \tilde{\mathbf{x}}_k + \tilde{g}_k \mathbb{E}\{\mathbf{x}_0|\tilde{\mathbf{x}}_k, \mathbf{y}\} + \tilde{\varsigma}_k \tilde{\mathbf{n}}_k \quad (44)$$

for

$$\tilde{h}_k \triangleq \sqrt{\frac{1 - \bar{\alpha}_{k-1} - \tilde{\varsigma}_k^2}{1 - \bar{\alpha}_k}} \quad \text{and} \quad \tilde{g}_k \triangleq \sqrt{\bar{\alpha}_{k-1}} - \tilde{h}_k \sqrt{\bar{\alpha}_k}. \quad (45)$$

Thus the VP DDIM reverse process can be described by (40), (44), and (45) with $\tilde{\mathbf{n}}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \forall k$ and initialization $\tilde{\mathbf{x}}_K \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

C DDIM details for VE

We now provide the details for the VE version of DDIM. Starting with the VP DDIM reverse process (44), we can divide both sides by $\sqrt{\bar{\alpha}_{k-1}}$ to get

$$\frac{\tilde{\mathbf{x}}_{k-1}}{\sqrt{\bar{\alpha}_{k-1}}} = \frac{\tilde{h}_k \sqrt{\bar{\alpha}_k}}{\sqrt{\bar{\alpha}_{k-1}}} \frac{\tilde{\mathbf{x}}_k}{\sqrt{\bar{\alpha}_k}} + \frac{\tilde{g}_k}{\sqrt{\bar{\alpha}_{k-1}}} \mathbb{E}\{\mathbf{x}_0|\tilde{\mathbf{x}}_k, \mathbf{y}\} + \frac{\tilde{\varsigma}_k}{\sqrt{\bar{\alpha}_{k-1}}} \tilde{\mathbf{n}}_k \quad (46)$$

and leveraging the VP-to-VE relationship (37) to write

$$\mathbf{x}_{k-1} = h_k \mathbf{x}_k + g_k \mathbb{E}\{\mathbf{x}_0|\mathbf{x}_k, \mathbf{y}\} + \varsigma_k \mathbf{n}_k \quad \text{with} \quad h_k = \frac{\tilde{h}_k \sqrt{\bar{\alpha}_k}}{\sqrt{\bar{\alpha}_{k-1}}}, \quad g_k = \frac{\tilde{g}_k}{\sqrt{\bar{\alpha}_{k-1}}}, \quad \varsigma_k = \frac{\tilde{\varsigma}_k}{\sqrt{\bar{\alpha}_{k-1}}} \quad (47)$$

with $\mathbf{n}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \forall k$ and initialization $\mathbf{x}_K \sim \mathcal{N}(\mathbf{0}, \sigma_K^2 \mathbf{I})$. Plugging \tilde{g}_k from (45) into (47), we find

$$g_k = \frac{\sqrt{\bar{\alpha}_{k-1}} - \tilde{h}_k \sqrt{\bar{\alpha}_k}}{\sqrt{\bar{\alpha}_{k-1}}} = 1 - h_k. \quad (48)$$

Then plugging (40) into (47), we find

$$\varsigma_k = \frac{\eta_{\text{ddim}}}{\sqrt{\bar{\alpha}_{k-1}}} \sqrt{\frac{1 - \bar{\alpha}_{k-1}}{1 - \bar{\alpha}_k}} \sqrt{1 - \frac{\bar{\alpha}_k}{\bar{\alpha}_{k-1}}} = \eta_{\text{ddim}} \sqrt{\frac{1}{\bar{\alpha}_{k-1}} \frac{1 - \bar{\alpha}_{k-1}}{1 - \bar{\alpha}_k} \left(1 - \frac{\bar{\alpha}_k}{\bar{\alpha}_{k-1}} \right)} \quad (49)$$

$$= \eta_{\text{ddim}} \sqrt{\frac{1 - \bar{\alpha}_{k-1}}{\bar{\alpha}_{k-1}} \frac{\bar{\alpha}_k}{1 - \bar{\alpha}_k} \left(\frac{1}{\bar{\alpha}_k} - \frac{1}{\bar{\alpha}_{k-1}} \right)} = \eta_{\text{ddim}} \sqrt{\frac{\sigma_{k-1}^2}{\sigma_k^2} ([1 + \sigma_k^2] - [1 + \sigma_{k-1}^2])} \quad (50)$$

$$= \eta_{\text{ddim}} \sqrt{\frac{\sigma_{k-1}^2 (\sigma_k^2 - \sigma_{k-1}^2)}{\sigma_k^2}}. \quad (51)$$

Finally, noting from (37), (47), and (51) that

$$\frac{\tilde{\varsigma}_k^2}{1 - \bar{\alpha}_{k-1}} = \varsigma_k^2 \frac{\bar{\alpha}_{k-1}}{1 - \bar{\alpha}_{k-1}} = \frac{\varsigma_k^2}{\sigma_{k-1}^2} = \frac{\eta_{\text{ddim}}^2 \sigma_{k-1}^2 (\sigma_k^2 - \sigma_{k-1}^2)}{\sigma_{k-1}^2 \sigma_k^2} = \eta_{\text{ddim}}^2 \left(1 - \frac{\sigma_{k-1}^2}{\sigma_k^2} \right), \quad (52)$$

we plug \tilde{h}_k from (45) into (47) to find

$$h_k = \sqrt{\frac{\bar{\alpha}_k}{\bar{\alpha}_{k-1}} \frac{1 - \bar{\alpha}_{k-1} - \tilde{\zeta}_k^2}{1 - \bar{\alpha}_k}} = \sqrt{\frac{\bar{\alpha}_k}{1 - \bar{\alpha}_k} \frac{1 - \bar{\alpha}_{k-1} - \tilde{\zeta}_k^2}{\bar{\alpha}_{k-1}}} \quad (53)$$

$$= \sqrt{\frac{\bar{\alpha}_k}{1 - \bar{\alpha}_k} \frac{1 - \bar{\alpha}_{k-1}}{\bar{\alpha}_{k-1}} \left(1 - \frac{\tilde{\zeta}_k^2}{1 - \bar{\alpha}_{k-1}}\right)} = \sqrt{\frac{\sigma_{k-1}^2}{\sigma_k^2} \left(1 - \eta_{\text{ddim}}^2 \left(1 - \frac{\sigma_{k-1}^2}{\sigma_k^2}\right)\right)} \quad (54)$$

$$= \sqrt{\frac{\sigma_{k-1}^2}{\sigma_k^2} \left(1 - \frac{\zeta_k^2}{\sigma_{k-1}^2}\right)} = \sqrt{\frac{\sigma_{k-1}^2}{\sigma_k^2} - \frac{\zeta_k^2}{\sigma_k^2}} = \sqrt{\frac{\sigma_{k-1}^2 - \zeta_k^2}{\sigma_k^2}}. \quad (55)$$

The VE DDIM reverse process is summarized in (26)-(27).

D Speeding up CG

In this section, we describe a small modification to FIRE that can help to speed up the CG step. When CG is used to solve (11), its convergence speed is determined by the condition number of $\mathbf{A}^\top \mathbf{A} + (\sigma_y^2/\nu)\mathbf{I}$ (Luenberger & Ye, 2016). Thus CG can converge slowly when σ_y^2/ν is small, which can happen in early DDfire iterations. To speed up CG, we propose to solve (11) using $\hat{\sigma}_y$ in place of σ_y , for some $\hat{\sigma}_y > \sigma_y$. Since the condition number of $\mathbf{A}^\top \mathbf{A} + (\hat{\sigma}_y^2/\nu)\mathbf{I}$ is at most $\nu s_{\max}^2/\hat{\sigma}_y^2 + 1$, we can guarantee a conditional number of at most 10001 by setting

$$\hat{\sigma}_y^2 = \nu s_{\max}^2 \max\{10^{-4}, \sigma_y^2/(\nu s_{\max}^2)\}. \quad (56)$$

Although using $\hat{\sigma}_y > \sigma_y$ in (11) will degrade the MSE of $\hat{\mathbf{x}}$, the degradation is partially offset by the fact that less noise will be added when renoising \mathbf{r} . In any case, the modified (11) can be written as

$$\hat{\mathbf{x}} = (\mathbf{A}^\top \mathbf{A}/\hat{\sigma}_y^2 + \mathbf{I}/\nu)^{-1} (\mathbf{A}^\top \mathbf{y}/\hat{\sigma}_y^2 + \bar{\mathbf{x}}/\nu) \quad (57)$$

$$= (\mathbf{A}^\top \mathbf{A}/\hat{\sigma}_y^2 + \mathbf{I}/\nu)^{-1} (\mathbf{A}^\top [\mathbf{A}\mathbf{x}_0 + \sigma_y \mathbf{w}]/\hat{\sigma}_y^2 + [\mathbf{x}_0 - \sqrt{\nu} \mathbf{e}]/\nu) \quad (58)$$

$$= \mathbf{x}_0 + (\mathbf{A}^\top \mathbf{A}/\hat{\sigma}_y^2 + \mathbf{I}/\nu)^{-1} (\mathbf{A}^\top \mathbf{w} \sigma_y / \hat{\sigma}_y^2 - \mathbf{e}/\sqrt{\nu}), \quad (59)$$

in which case $\hat{\mathbf{x}} \sim \mathcal{N}(\mathbf{x}_0, \mathbf{C})$ with covariance

$$\mathbf{C} = (\mathbf{A}^\top \mathbf{A}/\hat{\sigma}_y^2 + \mathbf{I}/\nu)^{-1} (\mathbf{A}^\top \mathbf{A} \sigma_y^2 / \hat{\sigma}_y^4 + \mathbf{I}/\nu) (\mathbf{A}^\top \mathbf{A}/\hat{\sigma}_y^2 + \mathbf{I}/\nu)^{-1} \quad (60)$$

$$= (\mathbf{V} \mathbf{S}^\top \mathbf{S} \mathbf{V}^\top / \hat{\sigma}_y^2 + \mathbf{I}/\nu)^{-1} (\mathbf{V} \mathbf{S}^\top \mathbf{S} \mathbf{V}^\top \sigma_y^2 / \hat{\sigma}_y^4 + \mathbf{I}/\nu) (\mathbf{V} \mathbf{S}^\top \mathbf{S} \mathbf{V}^\top / \hat{\sigma}_y^2 + \mathbf{I}/\nu)^{-1} \quad (61)$$

$$= \mathbf{V} \text{Diag}(\boldsymbol{\gamma}) \mathbf{V}^\top \text{ for } \gamma_i = \frac{s_i^2 \sigma_y^2 / \hat{\sigma}_y^4 + 1/\nu}{[s_i^2 / \hat{\sigma}_y^2 + 1/\nu]^2}. \quad (62)$$

The desired renoising variance then becomes

$$\boldsymbol{\Sigma} = \sigma^2 \mathbf{I} - \mathbf{C} = \mathbf{V} \text{Diag}(\boldsymbol{\lambda}) \mathbf{V}^\top \text{ for } \lambda_i = \sigma^2 - \frac{s_i^2 \sigma_y^2 / \hat{\sigma}_y^4 + 1/\nu}{[s_i^2 / \hat{\sigma}_y^2 + 1/\nu]^2} \quad (63)$$

and we can generate the colored noise \mathbf{c} via (16) if the SVD is practical. If not, we approximate $\boldsymbol{\Sigma}$ by

$$\hat{\boldsymbol{\Sigma}} = (\sigma^2 - \nu) \mathbf{I} + \xi \mathbf{A}^\top \mathbf{A} \text{ with } \xi = \frac{1}{s_{\max}^2} \left(\nu - \frac{s_i^2 \sigma_y^2 / \hat{\sigma}_y^4 + 1/\nu}{[s_i^2 / \hat{\sigma}_y^2 + 1/\nu]^2} \right) \quad (64)$$

and generate the colored noise \mathbf{c} via (19). It is straightforward to show that $\xi \geq 0$ whenever $\hat{\sigma}_y \geq \sigma_y$, in which case $\sigma^2 \geq \nu$ guarantees that $\hat{\boldsymbol{\Sigma}}$ is a valid covariance matrix. Figure 6 shows the close agreement between the ideal and approximate $\hat{\boldsymbol{\Sigma}}$ -renoisied error spectra both when $\hat{\sigma}_y = \sigma_y$ and when $\hat{\sigma}_y > \sigma_y$.




Figure 6: For FFHQ Gaussian deblurring, the left plot shows the eigenspectrum of the error covariance $\text{Cov}\{\bar{\mathbf{x}} - \mathbf{x}_0\}$ with either $\hat{\sigma}_y^2$ from (56) (if CG speedup) or $\hat{\sigma}_y^2 = \sigma_y^2$ (if no CG speedup), as well as the eigenspectrum of the target error covariance $\sigma^2 \mathbf{I}$ to aim for when renoising. The right plot shows the eigenvalues of the renoised error covariance $\text{Cov}\{\mathbf{r} - \mathbf{x}_0\}$ for the ideal case when Σ is used (possible with SVD) and the case when $\hat{\Sigma}$ from (17) is used (if no SVD), with either $\hat{\sigma}_y^2$ or σ_y^2 . Here we used $\sigma_y^2 = 10^{-6}$, $\nu = 0.16$ (corresponding to the first FIRE iteration of the first DDIM step), and $\rho = 35.7$ (corresponding to the example in Fig. 3).

E Proof of Theorem 1

To prove Theorem 1, we begin by writing the key FIRE steps with explicit iteration index $n \geq 1$:

$$\bar{\mathbf{x}}[n] = \mathbf{d}_\theta(\mathbf{r}[n], \sigma[n]) \quad (65)$$

$$\hat{\mathbf{x}}[n] = \left(\mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^2}{\nu[n]} \mathbf{I} \right)^{-1} \left(\mathbf{A}^\top \mathbf{y} + \frac{\sigma_y^2}{\nu[n]} \bar{\mathbf{x}}[n] \right) \quad (66)$$

$$\sigma^2[n+1] = \max\{\sigma^2[n]/\rho, \nu[n]\} \quad (67)$$

$$\lambda_i[n] = \sigma^2[n+1] - (s_i^2/\sigma_y^2 + 1/\nu[n])^{-1}, \quad i = 1, \dots, d \quad (68)$$

$$\mathbf{n}[n] = \mathbf{V} \text{Diag}(\boldsymbol{\lambda}[n])^{1/2} \boldsymbol{\epsilon}[n], \quad \boldsymbol{\epsilon}[n] \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (69)$$

$$\mathbf{r}[n+1] = \hat{\mathbf{x}}[n] + \mathbf{n}[n] \quad (70)$$

Our proof uses induction. By the assumptions of the theorem, we know that there exists an iteration n (in particular $n = 1$) for which $\mathbf{r}[n] = \mathbf{x}_0 + \sigma[n]\boldsymbol{\epsilon}[n]$ with $\boldsymbol{\epsilon}[n] \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and finite $\sigma[n]$. Then due to the denoiser assumption, we know that $\bar{\mathbf{x}}[n] = \mathbf{x}_0 - \sqrt{\nu[n]}\mathbf{e}[n]$ with $\mathbf{e}[n] \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and known $\nu[n] < \sigma^2[n]$. We assume that this value of $\nu[n]$ is used in lines (66)-(68). Using these results and (6), we can rewrite (66) as

$$\hat{\mathbf{x}}[n] = \left(\mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^2}{\nu[n]} \mathbf{I} \right)^{-1} \left(\mathbf{A}^\top (\mathbf{A}\mathbf{x}_0 + \sigma_y \mathbf{w}) + \frac{\sigma_y^2}{\nu[n]} (\mathbf{x}_0 - \sqrt{\nu[n]}\mathbf{e}[n]) \right) \quad (71)$$

$$= \mathbf{x}_0 + \left(\mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^2}{\nu[n]} \mathbf{I} \right)^{-1} \left(\sigma_y \mathbf{A}^\top \mathbf{w} - \frac{\sigma_y^2}{\sqrt{\nu[n]}} \mathbf{e}[n] \right) \sim \mathcal{N}(\mathbf{x}_0, \mathbf{C}[n]) \quad (72)$$

for

$$\mathbf{C}[n] \triangleq \left(\mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^2}{\nu[n]} \mathbf{I} \right)^{-1} \left(\sigma_y^2 \mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^4}{\nu[n]} \mathbf{I} \right) \left(\mathbf{A}^\top \mathbf{A} + \frac{\sigma_y^2}{\nu[n]} \mathbf{I} \right)^{-1} = \left(\frac{1}{\sigma_y^2} \mathbf{A}^\top \mathbf{A} + \frac{1}{\nu[n]} \mathbf{I} \right)^{-1} \quad (73)$$

by leveraging the independent-Gaussian assumption on $\mathbf{e}[n]$. From this and (68)-(69), we can then deduce

$$\mathbf{n}[n] \sim \mathcal{N}(\mathbf{0}, \Sigma[n]) \text{ with } \Sigma[n] \triangleq \mathbf{V} \text{Diag}(\boldsymbol{\lambda}[n]) \mathbf{V}^\top = \sigma^2[n+1] \mathbf{I} - \mathbf{C}[n] \quad (74)$$

Algorithm 3 FIRE for the GLM: $\hat{\mathbf{x}} = \text{FIRE}_{\text{GLM}}(\mathbf{y}, \mathbf{A}, p_{\mathbf{y}|\mathbf{z}}, \mathbf{r}_{\text{init}}, \sigma_{\text{init}}, N, \rho)$.

Require: $\mathbf{y}, \mathbf{A}, s_{\text{max}}, p_{\mathbf{y}|\mathbf{z}}, N, \rho > 1, \mathbf{r}_{\text{init}}, \sigma_{\text{init}}$. Also $\mathbf{A} = \mathbf{U} \text{Diag}(\mathbf{s}) \mathbf{V}^\top$ if using SVD.

```

1:  $\mathbf{r} = \mathbf{r}_{\text{init}}$  and  $\sigma = \sigma_{\text{init}}$  ▷ Initialize
2: for  $n = 1, \dots, N$  do
3:    $\bar{\mathbf{x}} \leftarrow \mathbf{d}_\theta(\mathbf{r}, \sigma) + \sqrt{\hat{\nu}_\phi(\sigma)} \mathbf{v}, \quad \mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  ▷ Stochastic denoising
4:    $\nu \leftarrow 2\hat{\nu}_\phi(\sigma)$  ▷ Error variance of  $\bar{\mathbf{x}}$ 
5:    $\bar{\mathbf{z}} \leftarrow \mathbf{A}\bar{\mathbf{x}}$  ▷ Denoised version of  $\mathbf{A}\mathbf{x}_0$ 
6:    $\bar{\nu}_z \leftarrow \nu \|\mathbf{A}\|_F^2 / m$  ▷ Error variance of  $\bar{\mathbf{z}}$ 
7:    $\hat{z}_j \leftarrow \mathbb{E}\{z_{0,j} | y_j; \bar{z}_j, \bar{\nu}_z\} \quad \forall j = 1, \dots, m$  ▷ Estimate  $z_{0,j} \sim \mathcal{N}(\bar{z}_j, \bar{\nu}_z)$  from  $y_j \sim p_{\mathbf{y}|\mathbf{z}}(\cdot | z_{0,j})$ 
8:    $\hat{\nu}_z \leftarrow \frac{1}{m} \sum_{j=1}^m \text{var}\{z_{0,j} | y_j; \bar{z}_j, \bar{\nu}_z\}$  ▷ Averaged posterior variance of  $\{z_{0,j}\}$ 
9:    $\bar{\sigma}_y^2 \leftarrow [1/\hat{\nu}_z - 1/\bar{\nu}_z]^{-1}$  ▷ Extrinsic variance
10:   $\bar{\mathbf{y}} \leftarrow \bar{\sigma}_y^2 (\hat{\mathbf{z}}/\hat{\nu}_z - \bar{\mathbf{z}}/\bar{\nu}_z)$  ▷ Extrinsic mean
11:   $\nu \leftarrow (\|\bar{\mathbf{y}} - \mathbf{A}\bar{\mathbf{x}}\|^2 - \bar{\sigma}_y^2 m) / \|\mathbf{A}\|_F^2$  ▷ Error variance of  $\bar{\mathbf{x}}$ 
12:   $\hat{\mathbf{x}} \leftarrow \arg \min_{\mathbf{x}} \|\bar{\mathbf{y}} - \mathbf{A}\mathbf{x}\|^2 / \bar{\sigma}_y^2 + \|\mathbf{x} - \bar{\mathbf{x}}\|^2 / \nu$  ▷ Estimate  $\mathbf{x}_0 \sim \mathcal{N}(\bar{\mathbf{x}}, \nu \mathbf{I})$  from  $\bar{\mathbf{y}} \sim \mathcal{N}(\mathbf{A}\mathbf{x}_0; \bar{\sigma}_y^2 \mathbf{I})$ 
13:   $\sigma^2 \leftarrow \max\{\sigma^2 / \rho, \nu\}$  ▷ Decrease target variance
14:  if have SVD then
15:     $\lambda_i \leftarrow \sigma^2 - (s_i^2 / \sigma_y^2 + 1/\nu)^{-1}, \quad i = 1, \dots, d$ 
16:     $\mathbf{c} \leftarrow \mathbf{V} \text{Diag}(\boldsymbol{\lambda})^{1/2} \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  ▷ Colored Gaussian noise
17:  else
18:     $\xi \leftarrow (\nu - (s_{\text{max}}^2 / \sigma_y^2 + 1/\nu)^{-1}) / s_{\text{max}}^2$ 
19:     $\mathbf{c} \leftarrow [\sqrt{\sigma^2 - \nu} \mathbf{I} \quad \sqrt{\xi} \mathbf{A}^\top] \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  ▷ Colored Gaussian noise
20:   $\mathbf{r} \leftarrow \hat{\mathbf{x}} + \mathbf{c}$  ▷ Renoise so that  $\text{Cov}\{\mathbf{r} - \mathbf{x}_0\} = \sigma^2 \mathbf{I}$ 
21: return  $\hat{\mathbf{x}}$ 

```

so that, from (70),

$$\mathbf{r}[n+1] \sim \mathcal{N}(\mathbf{x}_0, \mathbf{C}[n] + \boldsymbol{\Sigma}[n]) = \mathcal{N}(\mathbf{x}_0, \sigma^2[n+1] \mathbf{I}) \quad (75)$$

$$\Leftrightarrow \mathbf{r}[n+1] = \mathbf{x}_0 + \sigma[n+1] \boldsymbol{\varepsilon}[n+1], \quad \boldsymbol{\varepsilon}[n+1] \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (76)$$

Thus, by induction, if $\mathbf{r}[n] = \mathbf{x}_0 + \sigma[n] \boldsymbol{\varepsilon}[n]$ with $\boldsymbol{\varepsilon}[n] \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ holds at $n = 1$, then it holds at all $n > 1$.

Recall that the theorem also assumed that $\nu[n] < \sigma^2[n]$ for all n . Thus, there exists a $\rho > 1$ for which $\sigma^2[n]/\rho > \nu[n]$ for all n , for which we can rewrite (67) as

$$\sigma^2[n+1] = \sigma^2[n]/\rho \quad \forall n. \quad (77)$$

Consequently, for any iteration $n \geq 1$ we can write

$$\sigma^2[n] = \sigma^2[1]/\rho^{n-1} = \sigma_{\text{init}}^2/\rho^{n-1}. \quad (78)$$

Finally, because the error covariance on $\hat{\mathbf{x}}[n]$ obeys

$$\mathbf{C}[n] = \left(\frac{1}{\sigma_y^2} \mathbf{A}^\top \mathbf{A} + \frac{1}{\nu[n]} \mathbf{I} \right)^{-1} < \nu[n] \mathbf{I} < \sigma^2[n] \mathbf{I} = \frac{\sigma_{\text{init}}^2}{\rho^{n-1}} \mathbf{I} \quad (79)$$

we see that the error variance in $\hat{\mathbf{x}}[n]$ decreases exponentially with n and thus $\hat{\mathbf{x}}[n]$ converges to the true \mathbf{x}_0 .

F GLM-FIRE algorithm

The full algorithm for the GLM version of FIRE is given in Alg. 3. There, the dashed blue box indicates the lines used for the EP update; the remaining lines mirror those in SLM-FIRE, summarized as Alg. 1.

Table 5: Hyperparameter values used for DDfire.

Dataset	σ_y	Inpaint (box)		Deblur (Gaussian)		Deblur (Motion)		4× Super-resolution	
		K	δ	K	δ	K	δ	K	δ
FFHQ	0.05	100	0.50	650	0.60	500	0.20	650	0.60
ImageNet	0.05	100	0.50	500	0.20	500	0.20	650	0.60

G Implementation details

G.1 Inverse problems

For the linear inverse problems, the measurements were generated as

$$\mathbf{y} = \mathbf{A}\mathbf{x}_0 + \sigma_y \mathbf{w}, \quad \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (80)$$

with appropriate \mathbf{A} . For box inpainting, Gaussian deblurring, and super-resolution we used the \mathbf{A} and \mathbf{A}^\top implementations from Kavar et al. (2022b). For motion deblurring, we implemented our own \mathbf{A} and \mathbf{A}^\top with reflect padding. All methods used these operators implementations except DiffPIR, which used the authors’ implementations. Motion-blur kernels were generated using Borodenko (2020).

For phase retrieval, the measurements were generated using the method from Metzler et al. (2018):

$$y_j^2 = |z_{0,j}|^2 + w_j, \quad w_j \sim \mathcal{N}(0, \alpha_{\text{shot}}^2 |z_{0,j}|^2), \quad j = 1, \dots, m, \quad (81)$$

where α_{shot} controls the noise level and $\mathbf{z}_0 = \mathbf{A}\mathbf{x}_0$, with the values of \mathbf{x}_0 scaled to lie in the range $[0, 255]$. This is an approximation of the Poisson shot-noise corruption model in that the intensity $y_j^2/\alpha_{\text{shot}}^2$ is approximately $\text{Poisson}(|z_{0,j}|/\alpha_{\text{shot}})$ distributed for sufficiently small values of α_{shot} . We implemented the oversampled-Fourier \mathbf{A} by zero-padding the image by $2\times$ in each direction and then passing the result through a unitary FFT. For CDP phase retrieval, we set $\mathbf{A} = [\mathbf{A}_1^\top, \dots, \mathbf{A}_L^\top]^\top$ for $\mathbf{A}_l = L^{-1/2} \mathbf{F} \text{Diag}(\mathbf{c}_l)$, where \mathbf{F} is a $d \times d$ FFT and \mathbf{c}_l contain i.i.d. random entries uniformly distributed on the unit circle in the complex plane, and where $L = 4$. In both cases, $\mathbf{A}^\top \mathbf{A} = \mathbf{I}$.

G.2 Evaluation protocol

For the linear inverse problems, we run each method once for each measurement \mathbf{y} in the 1000-sample test set and compute average PSNR, average LPIPS, and FID from the resulting recoveries.

For OSF phase retrieval, following Chung et al. (2023a), we run each algorithm four times and keep the reconstruction $\hat{\mathbf{x}}$ that minimizes the measurement residual $\|\mathbf{y} - |\mathbf{A}\hat{\mathbf{x}}|\|$. Performance metrics are then evaluated after resolving the inherent spatial shift and conjugate flip ambiguities associated with phase retrieval (see, e.g., Bendory et al. (2015)). Note global phase ambiguity is not an issue due to the non-negativity of our images. For the CDP experiments, we run each algorithm only once and don’t perform ambiguity resolution, because it is unnecessary.

G.3 Unconditional diffusion models

For the FFHQ experiments, all methods used the pretrained model from Chung et al. (2023a). For the ImageNet experiments, all methods used the pretrained model from Dhariwal & Nichol (2021). In both cases, $T = 1000$.

G.4 Recovery methods

DDfire. Our Python/Pytorch codebase is a modification of the DPS codebase from Chung et al. (2023b), and our code will be released upon the acceptance of this paper. For all but one row of the ablation study

Table 6: Hyperparameter values used for DDfire phase retrieval.

Operator	α_{shot}	K	δ
OSF	8	20	0.00
CDP	45	80	0.00

in Table 1 and the dashed line in Fig. 4, we ran DDfire without an SVD and thus with the approximate renoising in (19).

For the linear inverse problems, unless noted otherwise, we ran DDfire for 1000 NFEs using $\eta_{\text{ddim}} = 1.5$, and we did not use stochastic denoising (i.e., $\hat{\nu}_\phi(\sigma) = 0 \forall \sigma$ in Alg. 1), as suggested by our ablation study. We tuned the (K, δ) hyperparameters to minimize LPIPS on a 100-sample validation set, yielding the parameters in Table 5. For the runtime results in Fig. 4, we used $\eta_{\text{ddim}} = 1.0$ for $N_{\text{tot}} \in \{50, 100, 200, 500\}$ and $\eta_{\text{ddim}} = 0$ for $N_{\text{tot}} = 20$, and we used $K = \text{NFE}/2$ and $\delta = 0.2$ for all cases. For the ν -estimation step in Alg. 1, we used $\|\mathbf{A}\|_F^2 \approx \frac{1}{L} \sum_{l=1}^L \|\mathbf{A}\mathbf{w}_l\|^2$ with i.i.d. $\mathbf{w}_l \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $L = 25$.

For phase retrieval, we ran DDfire for 800 NFEs by default and used $\eta_{\text{ddim}} = 0.85$ and the hand-tuned (K, δ) values listed in Table 6. Also, we did use stochastic denoising, since it helped in all metrics. Since the likelihood $p_{y|z}(y|z) = \mathcal{N}(y; |z|, \sigma_y^2)$ makes the conditional mean and variance in lines 8-7 of Alg. 3 intractable, we used the Laplace approximation (Bishop, 2007).

DDRM. We ran DDRM for 20 NFEs using the authors’ implementation from Kawar et al. (2022b) with minor changes to work with our codebase.

DiffPIR. We ran DiffPIR for 20 NFEs using the authors’ implementation from Zhu et al. (2024) without modification. Hyperparameters were set according to the reported values in Zhu et al. (2023).

IGDM. We ran IGDM for 100 NFEs. Since the authors do not provide a IGDM implementation for noisy inverse problems in NVlabs (2023), we coded IGDM ourselves in Python/PyTorch. With problems for which an SVD is available, we computed $(\mathbf{A}\mathbf{A}^\top + \zeta_k \mathbf{I})^{-1}$ using the efficient SVD implementation of \mathbf{A} from the DDRM codebase Kawar et al. (2022b), and otherwise we used CG.

DPS. For the linear inverse problems, we ran DPS for 1000 NFEs using the authors’ implementation from Chung et al. (2023b) without modification, using the suggested tuning from Chung et al. (2023a, Sec. D.1).

For phase retrieval, we also used 1000 NFEs, but made minor adjustments to the DPS authors’ implementation to accommodate the likelihood $p_{y|z}(y|z) = \mathcal{N}(y; |z|, \sigma_y^2)$, which was used by all methods for fairness. We used grid-search tuning to minimize LPIPS on a 100-image validation set.

RED-diff. We ran RED-diff for 1000 NFEs using the authors’ implementation from NVlabs (2023), with minor changes to work with our codebase. We tuned the RED-diff learning rate, λ , and data fidelity weight v_t to minimize LPIPS with a 100-image validation set. For phase retrieval, we made similar code modifications to handle our likelihood as with DPS.

DAPS. We ran DAPS for 1000 NFEs using the authors’ implementation from Zhang et al. (2024), with minor changes to work in our codebase. The tuning parameters were set as in Zhang et al. (2025). For phase retrieval, we made similar code modifications to handle our likelihood as with DPS.

DOLPH. We ran DOLPH for 1000 NFEs. Since the DOLPH authors did not release an implementation, we implemented it ourselves in Python/PyTorch and used grid-search to find the step-size that minimized LPIPS on a 100-image validation set. For phase retrieval, we made similar code modifications to handle our likelihood as with DPS.




Figure 7: LPIPS vs. single image sampling time for noisy OSF phase retrieval on an A100 GPU. The evaluation used 1000 FFHQ images.

HIO. We translated the MATLAB implementation of HIO from Metzler (2018) to Python and set the step-size parameter to 0.9. We then followed the runtime procedure described in Metzler et al. (2018): For the OSF experiments, HIO is first run 50 times, for 50 iterations each, from a random initialization. The estimate $\hat{\mathbf{x}}$ with the lowest measurement residual $\|\mathbf{y} - |\mathbf{A}\hat{\mathbf{x}}|\|$ is then used to reinitialize HIO, after which it is run for 1000 more iterations. Finally, the second and third color channels in the result are shifted and flipped as needed to best match the first color channel. For the CDP experiments, HIO is run once for 200 iterations from a random initialization.

prDeep. We used the Python implementation from Hekstra et al. (2018). As recommended in Metzler et al. (2018), we initialized prDeep with the HIO estimate for OSF experiments and with an all-ones initialization for CDP experiments. (Note that only prDeep uses the HIO initialization; DPS, DOLPH, and DDfire do not.) We tuned λ on a grid to minimize LPIPS on a 100-image validation set.

G.5 Compute

All experiments were run on a single NVIDIA A100 GPU with 80GB of memory. The runtime for each method on the GPU varies, as shown in Figure 4.

H Additional experimental results

Figure 7 shows LPIPS vs. average runtime (in seconds on an A100 GPU) to generate a single image for noisy OSF phase retrieval on the 1000-sample FFHQ test set. The figure shows that DDfire gives a significantly better performance/complexity tradeoff than all diffusion-based competitors.

Figure 8 shows SLM-FIRE’s σ^2 versus iteration i , for comparison to the true denoiser input variance $\|\mathbf{r} - \mathbf{x}_0\|_2^2/d$, and SLM-FIRE’s ν , for comparison to the true denoiser output variance $\|\bar{\mathbf{x}}_0 - \mathbf{x}_0\|_2^2/d$, for 25 FIRE iterations with $\rho = 1.5$ for noisy $4\times$ super-resolution at $t[k] = 1000$. We see that the SLM-FIRE estimates σ^2 and ν track the true error variances quite closely.

Figure 9 shows a similar figure for GLM-FIRE. In particular, it shows GLM-FIRE’s σ^2 versus iteration i , for comparison to the true denoiser input variance $\|\mathbf{r} - \mathbf{x}_0\|_2^2/d$, GLM-FIRE’s ν , for comparison to the true denoiser output variance $\|\bar{\mathbf{x}}_0 - \mathbf{x}_0\|_2^2/d$, and GLM-FIRE’s $\bar{\sigma}_y^2$, for comparison to the true pseudo-measurement variance $\|\bar{\mathbf{y}} - \mathbf{A}\mathbf{x}_0\|^2/m$ for noisy FFHQ phase retrieval at $t[k] = 1000$. We see that the GLM-FIRE estimates σ^2 , ν , and $\bar{\sigma}_y^2$ track the true variances quite closely.

Example recoveries for noisy phase retrieval with FFHQ images are given in Figure 10.

Additional example recoveries for the noisy linear inverse problems with FFHQ images are shown in Figure 11.




Figure 8: SLM-FIRE σ^2 , true denoiser input variance $\|\mathbf{r} - \mathbf{x}_0\|_2^2/d$, SLM-FIRE ν , and true denoiser output variance $\|\bar{\mathbf{x}}_0 - \mathbf{x}_0\|_2^2/d$ vs. SLM-FIRE iteration for noisy $4\times$ super-resolution at $t[k] = 1000$ for a single validation sample \mathbf{x}_0 .




Figure 9: GLM-FIRE σ^2 , true denoiser input variance $\|\mathbf{r} - \mathbf{x}_0\|_2^2/d$, GLM-FIRE ν , and true denoiser output variance $\|\bar{\mathbf{x}}_0 - \mathbf{x}_0\|_2^2/d$, GLM-FIRE $\bar{\sigma}_y^2$, and noise variance $\|\bar{\mathbf{y}}_0 - \mathbf{A}\mathbf{x}_0\|_2^2/m$ vs. GLM-FIRE iteration for noisy CDP phase retrieval at $t[k] = 1000$ for a single validation sample \mathbf{x}_0 .



Figure 10: Example recoveries from noisy phase retrieval with FFHQ images.



Figure 11: Example recoveries from noisy linear inverse problems with FFHQ images.