

Image Restoration by Denoising Diffusion Models with Iteratively Preconditioned Guidance

Tomer Garber
 Open University of Israel
 tomergarber@gmail.com

Tom Tirer
 Bar-Ilan University, Israel
 tirer.tom@gmail.com

Abstract

Training deep neural networks has become a common approach for addressing image restoration problems. An alternative for training a “task-specific” network for each observation model is to use pretrained deep denoisers for imposing only the signal’s prior within iterative algorithms, without additional training. Recently, a sampling-based variant of this approach has become popular with the rise of diffusion/score-based generative models. Using denoisers for general purpose restoration requires guiding the iterations to ensure agreement of the signal with the observations. In low-noise settings, guidance that is based on back-projection (BP) has been shown to be a promising strategy (used recently also under the names “pseudoinverse” or “range/null-space” guidance). However, the presence of noise in the observations hinders the gains from this approach. In this paper, we propose a novel guidance technique, based on preconditioning that allows traversing from BP-based guidance to least squares based guidance along the restoration scheme. The proposed approach is robust to noise while still having much simpler implementation than alternative methods (e.g., it does not require SVD or a large number of iterations). We use it within both an optimization scheme and a sampling-based scheme, and demonstrate its advantages over existing methods for image deblurring and super-resolution.

1. Introduction

Image restoration problems appear in a wide range of applications, where the goal is to recover a high-quality image $\mathbf{x}^* \in \mathbb{R}^n$ from its degraded version $\mathbf{y} \in \mathbb{R}^m$, which can be noisy, blurry, low-resolution, etc. In many problems, the relation between \mathbf{y} and \mathbf{x}^* can be expressed using a linear observation model

$$\mathbf{y} = \mathbf{A}\mathbf{x}^* + \mathbf{e} \quad (1)$$



Figure 1. Motion deblurring with noise level 0.05. From top to bottom and left to right: original, observation, DPS [6], DiffPIR [45] and our DDPG.

where $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a measurement operator with $m \leq n$, and $\mathbf{e} \sim \mathcal{N}(\mathbf{0}, \sigma_e^2 \mathbf{I}_m)$ is an additive white Gaussian noise. For example, in the denoising task \mathbf{A} is the identity matrix \mathbf{I}_n ; in the deblurring task \mathbf{A} is a blur operator; and in the super-resolution task \mathbf{A} is a composition of sub-sampling and (anti-aliasing) filtering. Image restoration problems are ill-posed: just looking for $\mathbf{x} \in \mathbb{R}^n$ that fits \mathbf{y} based on the observation model does not suffice for a successful recovery. Thus, it is required to utilize prior knowledge on the nature of \mathbf{x}^* .

Nowadays, it has become common to address these problems by exhaustively training a different deep neural network (DNN) for each *predefined* observation model in a supervised manner. Namely, using (1) to generate training pairs $\{\mathbf{y}_i, \mathbf{x}_i^*\}$ and training a DNN to invert the map [8, 18, 34, 43]. However, these “task-specific” DNNs suffer from a huge performance drop when the observations at test-time mismatch (even slightly) the assumptions made in training [1, 28, 37], which limits their applicability in many practical cases.

An alternative approach is to use pretrained DNNs that impose only the signal prior, while the agreement with the

observations is being handled at test-time in a “zero-shot” manner. A successful choice of such pretrained DNNs are Gaussian denoisers, as was initially demonstrated in different “plug-and-play” (PnP) and “regularization by denoising” (RED) iterative schemes [25, 36, 40, 44]. The rise of score-based generative modeling [13, 31] — typically referred to as denoising diffusion models (DDMs) — whose inference/sampling “reversed” flows are based on iteratively injecting Gaussian noise and denoising, has further increased the popularity of using iterative denoising for general purpose restoration. In particular, operations that promote data-fidelity, similar to those used in PnP/RED, have been used within the diffusion models’ iterative sampling schemes to guide the iterations towards an image that not only looks natural but also agrees with the observations [2, 5, 6, 16, 21, 30, 32, 41, 45].

Combining iterative denoising with guidance that is based on back-projection (BP) of intermediate estimates on the subspace $\{\mathbf{x} : \mathbf{A}\mathbf{x} = \mathbf{y}\}$ (more details in Section 2) was originally proposed in [36] (in the context of PnP), and has been shown to be a promising strategy in both the PnP and the DDM literature [5, 11, 19, 26, 27, 30, 37, 41, 46]. In the context of diffusion models it has been rediscovered under the names “pseudoinverse” [30] or “range/null-space” [41] guidance. However, in cases where \mathbf{A} has low singular values (e.g., in image deblurring and super-resolution), the presence of noise in the observations \mathbf{y} hinders the application of this approach: careful regularization is required and the performance gains tend to decrease.

Contribution. In this paper, we identify the BP step as a preconditioned version of a least squares (LS) gradient step. This inspires us to propose an iteration-dependent preconditioned guidance approach that essentially traverses from the BP step to the LS gradient step along the restoration scheme. Thus, it enjoys benefits of BP step [38, 39], such as stronger consistency with \mathbf{y} and accelerated convergence in early iterations, and the better robustness to noise of LS as the scheme approaches its end. The specific trajectory that we devise in this paper is also more simple and efficient to implement than other potential alternatives for BP, which either require having full knowledge on the SVD of \mathbf{A} [16], or many neural function evaluations (NFEs) [6].

We present formal mathematical motivation for our guidance technique and use it within both a PnP-based scheme and its DDM sampling-based counterpart that we design. We examine the proposed approach for image deblurring and super-resolution on the CelebA-HQ and ImageNet datasets, where we exploit the powerful denoisers of DDMs [7, 20], and demonstrate its advantages over existing methods. In particular, our sampling-based reconstruction approach is flexible to the observation model, computationally convenient, and displays a good combination of perceptual quality and accuracy.

2. Background and Related Work

2.1. PnP Denoisers and Back-Projections

Traditionally, estimating \mathbf{x}^* from its measurements (1) has been tackled by minimizing an explicit objective function

$$L(\mathbf{x}) = \ell(\mathbf{x}; \mathbf{y}) + s(\mathbf{x}), \quad (2)$$

which is composed of a data-fidelity term $\ell(\mathbf{x}; \mathbf{y})$ and a signal prior term $s(\mathbf{x})$. The PnP concept [40] suggests applying a proximal algorithm to (2), where the proximal mapping associated with the prior function $s(\cdot)$, i.e., $\text{prox}_s(\mathbf{x}) := \text{argmin}_{\mathbf{z}} \frac{1}{2} \|\mathbf{z} - \mathbf{x}\|_2^2 + s(\mathbf{z})$ [22], is replaced with an off-the-shelf Gaussian denoiser $\mathcal{D}(\cdot; \sigma_t)$ (σ_t denotes the denoiser’s noise level) that need not be associated with an explicit prior function (e.g., a pretrained DNN). A vast amount of literature has been dedicated to designing and analyzing PnP algorithms [4, 14, 24, 25, 35, 36, 40, 44].

A popular PnP scheme based on the proximal-gradient method (PGM) is given by:

$$\mathbf{x}_{0|t} = \mathcal{D}(\mathbf{x}_t; \sigma_t), \quad (3)$$

$$\mathbf{x}_{t-1} = \mathbf{x}_{0|t} - \mu_t \nabla_{\mathbf{x}} \ell(\mathbf{x}_{0|t}; \mathbf{y}), \quad (4)$$

where $t = T, T-1, \dots, 1$ is the iteration index (decreases over time, to match the notation in diffusion/score-based models), μ_t is a step-size, and $\sigma_T > \dots > \sigma_1$ is a predefined decreasing set of noise levels. The scheme is initialized with some \mathbf{x}_T and the final estimate is given by $\mathbf{x}_{0|1}$.

A typical choice of data-fidelity term is the LS objective

$$\ell_{LS}(\mathbf{x}; \mathbf{y}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2, \quad (5)$$

for which the gradient step in (4) is given by

$$\mathbf{x}_{t-1} = \mathbf{x}_{0|t} - \mu_t \mathbf{A}^T (\mathbf{A}\mathbf{x}_{0|t} - \mathbf{y}). \quad (6)$$

Note though that PnP methods based on LS gradient steps tend to require many iterations [25, 39]. The IDBP method [36] suggests that the data-fidelity guidance, which follows the denoising operation, will be the back-projection (BP) of $\mathbf{x}_{0|t}$ onto the affine subspace $\{\mathbf{x} : \mathbf{A}\mathbf{x} = \mathbf{y}\}$:

$$\mathbf{x}_{t-1} = \text{argmin}_{\mathbf{x}} \|\mathbf{x} - \mathbf{x}_{0|t}\|_2 \text{ s.t. } \mathbf{A}\mathbf{x} = \mathbf{y} \quad (7)$$

$$= \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I}_n - \mathbf{A}^\dagger \mathbf{A}) \mathbf{x}_{0|t} \quad (8)$$

$$= \mathbf{x}_{0|t} - \mathbf{A}^\dagger (\mathbf{A}\mathbf{x}_{0|t} - \mathbf{y}), \quad (9)$$

where \mathbf{A}^\dagger denotes the pseudoinverse of \mathbf{A} .

When \mathbf{A} has near-zero singular values, it has been shown to be beneficial to replace $\mathbf{A}^\dagger = \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^\dagger$ with a regularized version $\mathbf{A}^T (\mathbf{A}\mathbf{A}^T + \eta \mathbf{I}_m)^{-1}$, where $\eta > 0$ is the regularization hyperparameter [36, 37]. Importantly, note that there are popular tasks where the pseudoinverse can be

implemented efficiently (e.g., see the supplementary material for FFT based implementations of this operation for deblurring and super-resolution). And, in general, full rank $\mathbf{A}\mathbf{A}^T$ (and $\mathbf{A}\mathbf{A}^T + \eta\mathbf{I}_m$ otherwise) can be typically “inverted” using few conjugate gradient iterations [12], which only require applying \mathbf{A} and \mathbf{A}^T and bypass the need of matrix inversion or SVD.

As pointed out in [38, 39], (9) is equivalent to a gradient step (4) with step-size $\mu_t = 1$ and a special choice of $\ell(\mathbf{x}; \mathbf{y})$, which they dubbed as the “BP term”:

$$\ell_{BP}(\mathbf{x}; \mathbf{y}) = \frac{1}{2} \|(\mathbf{A}\mathbf{A}^T)^{-1/2}(\mathbf{A}\mathbf{x} - \mathbf{y})\|_2^2. \quad (10)$$

These works have analyzed, both empirically and theoretically, the effects of using BP steps (9) rather than LS steps (6) in inverse problems. They show provable faster convergence benefits (less iterations) and, in the low-noise regime, also MSE benefits. However, it is also shown that in the presence of observation noise and \mathbf{A} with small singular values, we have that $\mathbf{A}^\dagger \mathbf{y}$ amplifies the noise. Thus a strong regularization η is needed which hinders the advantages.

2.2. Restoration via Guidance of DDMs

Pretrained diffusion/score-based generative models [13, 31, 33] have been shown to be a powerful signal prior for image restoration. These generative models are based on training a network that approximates the score function $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$, where p_t denotes the distribution of \mathbf{x}_t , a Gaussian noisy version of the data $\mathbf{x}_0 \sim p_0$ with noise level indexed by $t \in [0, T]$. For $\mathbf{x}_t = \mathbf{x}_0 + \boldsymbol{\epsilon}_t$ with $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \sigma_t^2 \mathbf{I}_n)$, by Tweedie’s formula $-\sigma_t^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) = \mathbf{x}_t - \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \mathbb{E}[\boldsymbol{\epsilon}_t | \mathbf{x}_t]$ [9]. Therefore, these approaches essentially attempt to learn the MMSE Gaussian denoiser or, equivalently, noise estimation. After training, the sampling schemes are based on random Gaussian initialization and alternatively denoising and injection of Gaussian noise with a decreasing noise level.

For concreteness, let us focus on the popular DDPM formulation [13]. In DDPM, \mathbf{x}_t is generated from a data point \mathbf{x}_0 by the “forward flow”:

$$\mathbf{x}_t = \sqrt{1 - \beta_t} \mathbf{x}_{t-1} + \sqrt{\beta_t} \boldsymbol{\epsilon}, \quad (11)$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ and $\{\beta_t\}_{t=1}^T$ is a predetermined “noise schedule” obeying $0 < \beta_1 < \dots < \beta_T \leq 1$. By properties of Gaussians, \mathbf{x}_t can be directly generated from \mathbf{x}_0 by

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}, \quad (12)$$

which includes noise variance $1 - \bar{\alpha}_t$ and signal scale factor $\sqrt{\bar{\alpha}_t}$, where $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ and $\alpha_t = 1 - \beta_t$. In training time, a U-net DNN $\boldsymbol{\epsilon}_\theta(\cdot; t)$ is trained to predict the noise $\boldsymbol{\epsilon}$ given \mathbf{x}_t by minimizing an MSE loss. In DDPM’s inference time, starting from a random $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, one step of

the “reverse flow” that is used for generating a sample from p_0 is given by

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t; t) \right) + \sqrt{\beta_t} \boldsymbol{\epsilon}_t, \quad (13)$$

where $\boldsymbol{\epsilon}_t$ is drawn from $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$.

A faster generation scheme, denoted DDIM [29], is based on diverging from the Markovian guideline of DDPM. In DDIM, the step (13) is replaced with

$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \mathbf{x}_{0|t} + \hat{\sigma}_t \boldsymbol{\epsilon}_\theta(\mathbf{x}_t; t) + \tilde{\sigma}_t \boldsymbol{\epsilon}_t, \quad (14)$$

where $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ and

$$\hat{\sigma}_t = \sqrt{1 - \bar{\alpha}_{t-1} - \tilde{\sigma}_t^2}, \quad (15)$$

$$\mathbf{x}_{0|t} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t; t) \right). \quad (16)$$

The hyperparameter $\tilde{\sigma}_t$ allows trading between the levels of two noise terms: the estimated noise $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t; t)$ and a pure stochastic noise $\boldsymbol{\epsilon}_t$. Moreover, $\mathbf{x}_{0|t}$ in (16) is an estimate of \mathbf{x}_0 given \mathbf{x}_t , which is essentially performing Gaussian denoising: $\mathbf{x}_{0|t} = \frac{1}{\sqrt{\bar{\alpha}_t}} \mathcal{D}(\mathbf{x}_t; \sigma_t = \sqrt{1 - \bar{\alpha}_t})$, as implied by (12).

Many of the methods that use DDMs as priors for image restoration [16, 30, 41, 45] are based on the DDIM scheme [29]. Yet, in order that the sampling scheme will produce an image that is not only perceptually pleasing but also agrees with the measurements \mathbf{y} and the observation model (1) it is required to equip the iterations with some data-fidelity guidance. This guidance is typically based on the gradient of a data-fidelity term $\ell(\mathbf{x}; \mathbf{y})$. In other words, the iterations (14) are modified into

$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \mathbf{x}_{0|t} - \mu_t \nabla_{\mathbf{x}} \ell(\mathbf{x}_{0|t}; \mathbf{y}) + \hat{\sigma}_t \boldsymbol{\epsilon}_\theta(\mathbf{x}_t; t) + \tilde{\sigma}_t \boldsymbol{\epsilon}_t, \quad (17)$$

where μ_t is the guidance scaling factor.

Note the similarity between the PnP restoration (3)-(4) and the DDM restoration (16)-(17). Namely, ignoring the scaling of the signal by $\{\bar{\alpha}_t\}$, the main difference is the noise injection, composed of the estimated noise $\boldsymbol{\epsilon}_\theta$ and the random noise $\boldsymbol{\epsilon}_t$. The similarity of DDMs schemes to previous iterative denoising methods is further seen by the formulation of [15], which decouples the diffusion “forward” flows, used for training the denoisers, from the “reverse” inference flows. More discussion on such connections can be found in a recent review paper [10].

For example, the DDM restoration method in [6], dubbed DPS, where the guidance is based on the gradient of the LS term, resembles (6) with noise injection. Interestingly, just like its PnP counterpart, it requires a very large number of iterations (and thus many NFEs). Alternatively, by using BP guidance in (17) — equivalently, when injecting noise

to step (9) in IDBP — one gets the recent DDNM method [41] (which is also similar to the method in [30]). However, note that the DDNM paper focuses on noiseless settings and has difficulties in noisy settings (in Section 4 we state its technical limitation), which is aligned with the sensitivity of BP guidance to noise in \mathbf{y} , as mentioned above.

At this point, we would also like to mention the DDRM method [16], which resembles the BP approach as it generates “spectral space” measurements $\bar{\mathbf{y}}$ in a way similar to applying \mathbf{A}^\dagger on \mathbf{y} . Yet, it is more robust to observation noise as it requires access to the full SVD of \mathbf{A} , which allows it to mitigate noise amplification per singular component in each iteration. In this paper we aim to devise an approach that does not require computing and storing the SVD of \mathbf{A} , which is impractical in most cases.

Finally, we note that the way we construct a sampling scheme from a deterministic algorithm shares similarity with the recent method DiffPIR [45]. However, while [45] uses an existing PnP baseline, here we propose a novel core reconstruction method. Also, empirically our new sampling scheme demonstrates better accuracy than DiffPIR (evaluated by PSNR) without compromising on perceptual quality (evaluated by LPIPS).

3. The Proposed Method

In this section, we present our image restoration method that is more robust to observation noise than methods that use only back-projections while still having low computational complexity and clear benefits over methods that are based only on plain LS based guidance. To this end, we present a novel core guidance approach, equip it with formal theoretical motivation, and utilize it to devise a new sampling-based reconstruction scheme.

All the claims in this section are proved in the supplementary material.

3.1. Core approach

The idea for our approach comes from identifying (9) as a specific instance of the formula

$$\mathbf{x}_{t-1} = \mathbf{x}_{0|t} - \mu_t \mathbf{A}^T \mathbf{W} (\mathbf{A} \mathbf{x}_{0|t} - \mathbf{y}) \quad (18)$$

with $\mathbf{W} = (\mathbf{A} \mathbf{A}^T)^{-1}$ (assuming full rank \mathbf{A}) and $\mu_t = 1$. The plain LS step (6) is obviously also an instance of this expression with $\mathbf{W} = \mathbf{I}_m$.

Recall traditional preconditioning in optimization [23]: The optimization of an objective, e.g., $L : \mathbb{R}^n \rightarrow \mathbb{R}$ in (2), is based on gradient-based optimizers where the full gradient $\nabla L(\mathbf{x})$ is multiplied by a (symmetric) positive definite matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$. Such practice does not affect the problem’s minimizers (since $\mathbf{P} \nabla L(\mathbf{x}) = \mathbf{0} \iff \nabla L(\mathbf{x}) = \mathbf{0}$ due to the invertibility of \mathbf{P}), and its goal is to ease the optimization landscape, characterized by the condition number

of the effective Hessian $\mathbf{P}^{1/2} \nabla^2 L(\mathbf{x}) \mathbf{P}^{1/2}$. In (18), however, $\mathbf{W} \in \mathbb{R}^{m \times m}$ does not directly affect the signal’s prior and for general \mathbf{W} we may not have $\mathbf{A}^T \mathbf{W} (\mathbf{A} \mathbf{x} - \mathbf{y}) = \mathbf{0} \iff \nabla \ell_{LS}(\mathbf{x}) = \mathbf{A}^T (\mathbf{A} \mathbf{x} - \mathbf{y}) = \mathbf{0}$.

Nevertheless, we have the following claims.

Claim 3.1. *Let $\mathbf{A} \in \mathbb{A}^{m \times n}$ with $m \leq n$ and $\text{rank}(\mathbf{A}) = m$. Let $\mathbf{W} \in \mathbb{R}^{m \times m}$ such that $\text{rank}(\mathbf{A}^T \mathbf{W}) = m$. Then,*

$$\mathbf{A}^T \mathbf{W} (\mathbf{A} \mathbf{x} - \mathbf{y}) = \mathbf{0} \iff \mathbf{A}^T (\mathbf{A} \mathbf{x} - \mathbf{y}) = \mathbf{0}.$$

Claim 3.2. *Let $\mathbf{A} \in \mathbb{A}^{m \times n}$ with $m \leq n$. Let $\mathbf{W} \in \mathbb{R}^{m \times m}$ be a positive definite matrix that shares eigenbasis with $\mathbf{A} \mathbf{A}^T$. Then, there exists a positive definite $\mathbf{P} \in \mathbb{R}^{n \times n}$ such that*

$$\mathbf{A}^T \mathbf{W} \mathbf{A} = \mathbf{P}^{1/2} \mathbf{A}^T \mathbf{A} \mathbf{P}^{1/2}.$$

Therefore, if \mathbf{A} is full rank, then any \mathbf{W} for which $\mathbf{A}^T \mathbf{W}$ is full rank (e.g., $\mathbf{W} = (\mathbf{A} \mathbf{A}^T)^{-1}$) can be interpreted as a preconditioner of $\nabla \ell_{LS}$ (rather than of the full objective that may include additional terms). Alternatively, even if \mathbf{A} is not full rank, but \mathbf{W} shares eigenbasis with $\mathbf{A} \mathbf{A}^T$ (e.g., $\mathbf{W} = (\mathbf{A} \mathbf{A}^T + \eta \mathbf{I}_m)^{-1}$), we allow ourselves to refer to \mathbf{W} as a preconditioner of $\nabla \ell_{LS}$ since it has the same effect on the effective Hessian, $\mathbf{P}^{1/2} \nabla^2 \ell_{LS}(\mathbf{x}) \mathbf{P}^{1/2}$, as a traditional preconditioner \mathbf{P} .

Inspired by this, in lieu of (4) we propose a data-fidelity guidance based on *iteration-dependent* preconditioner:

$$\mathbf{x}_{t-1} = \mathbf{x}_{0|t} - \mu_t \mathbf{A}^T \mathbf{W}_t (\mathbf{A} \mathbf{x}_{0|t} - \mathbf{y}) \quad (19)$$

with a “smooth” discretization of the path of symmetric preconditioners $\{\mathbf{W}_t\}$ — invertible with eigenbasis shared with $\mathbf{A} \mathbf{A}^T$ — that starts roughly at $\mathbf{W}_T \propto (\mathbf{A} \mathbf{A}^T + \eta \mathbf{I}_m)^{-1}$ and ends roughly at $\mathbf{W}_0 \propto \mathbf{I}_m$.

In words, the proposed guidance is based on gently traversing from the (regularized) BP step to the LS step along the restoration scheme. Thus, it enjoys benefits of BP steps (stronger consistency with \mathbf{y} , accelerated convergence, etc.) in early iterations and the better robustness to noise of LS steps as the scheme approaches its end.

Potentially, a sequence of preconditioners $\{\mathbf{W}_t\}$ might be learned in a data-driven manner under some optimality criterion. Yet, this is expected to come at the price of losing their generality to arbitrary data. Thus, here we propose an engineered, but simple and effective, update rule of \mathbf{W}_t that, as we will show empirically, already demonstrates the usefulness of our new concept. Specifically, the rule that we propose is given by the following convex combination:

$$\mathbf{W}_t = (1 - \delta_t) (\mathbf{A} \mathbf{A}^T + \eta \mathbf{I}_m)^{-1} + \delta_t c \mathbf{I}_m \quad (20)$$

where $0 \leq \delta_T < \dots < \delta_0 \leq 1$ is a predefined sequence, and c is a positive scalar that balances the two terms.

In general, the main motivation for (20) is that it both obeys the previously stated conditions and also yields a

data-fidelity guidance that is very simple to implement. Indeed, plugging (20) in (19) yields

$$\begin{aligned} \mathbf{x}_{t-1} &= \mathbf{x}_{0|t} - \mu_t \left((1 - \delta_t) \mathbf{g}_{BP}(\mathbf{x}_{0|t}) + \delta_t \mathbf{g}_{LS}(\mathbf{x}_{0|t}) \right), \\ &=: \mathbf{x}_{0|t} - \mu_t \mathbf{g}_{\delta_t}(\mathbf{x}_{0|t}) \end{aligned} \quad (21)$$

where

$$\mathbf{g}_{BP}(\mathbf{x}_{0|t}) = \mathbf{A}^T (\mathbf{A} \mathbf{A}^T + \eta \mathbf{I}_m)^{-1} (\mathbf{A} \mathbf{x}_{0|t} - \mathbf{y}), \quad (22)$$

$$\mathbf{g}_{LS}(\mathbf{x}_{0|t}) = c \mathbf{A}^T (\mathbf{A} \mathbf{x}_{0|t} - \mathbf{y}), \quad (23)$$

can be computed efficiently for a wide range of observation models.

In Section 3.4 we will discuss practical implementations details for simple hyperparameter tuning (e.g., a simple way to tune $\{\delta_t\}$). As for setting c , observe that (19) is essentially a gradient step (at $\mathbf{x}_{0|t}$) of the weighted least squares (WLS) data-fidelity term:

$$\ell_{WLS,t}(\mathbf{x}; \mathbf{y}) = \frac{1}{2} \|\mathbf{W}_t^{1/2} (\mathbf{A} \mathbf{x} - \mathbf{y})\|_2^2, \quad (24)$$

for which we have the following property.

Claim 3.3. Denote by λ_1 the largest singular value of \mathbf{A} . Let $c \leq 1/\lambda_1^2$. Then, the update (21) with $\mu_t = 1$ ensures reduction in (24).

Many popular degradation models obey $\lambda_1 \lesssim 1$ (e.g., in deblurring $\lambda_1 = 1$ and in super-resolution λ_1 is moderately lower than 1). Thus, following the claim, unless stated otherwise, we use $c = 1$ and $\mu_t = 1$.

To conclude this subsection, the core of the proposed algorithm is based on iteratively computing (3) and (21). Without noise injection, this is essentially a novel PnP scheme, which we call Iterative Denoising and Preconditioned Guidance (IDPG). As we will empirically show in Section 4, while this method has good PSNR performance, turning it into a DDM-based sampling scheme leads to better recoveries, especially in terms of perceptual quality. Before presenting the sampling scheme in Section 3.3, let us provide some theoretical motivation for our preconditioned guidance approach.

3.2. Theory

Let us state some mathematical claims that motivate the core guidance approach. In this subsection we consider a simplified setting of having a fixed $0 < \delta < 1$, i.e., a fixed WLS data-fidelity term (24) with \mathbf{W} given as in (20) but with $\delta_t = \delta$. We examine the effects of this WLS, compared to LS and BP fidelity terms, on an estimator of \mathbf{x}^* , denoted by $\hat{\mathbf{x}}$, which is obtained by minimizing (2) with Tikhonov prior $s(\mathbf{x}) = \frac{\beta}{2} \|\mathbf{D} \mathbf{x}\|_2^2$, where $\beta > 0$ and $\mathbf{D}^T \mathbf{D}$ is invertible, so, the estimator $\hat{\mathbf{x}}$ per data-fidelity term is unique.

The MSE of $\hat{\mathbf{x}}$ (conditioned on \mathbf{x}^* , averaged over the noise \mathbf{e}) can be decomposed into two terms: squared bias, b^2 (independent of \mathbf{e}), and variance, v (depends on σ_e^2), given by

$$b^2 := \|\mathbb{E}[\hat{\mathbf{x}}|\mathbf{x}^*] - \mathbf{x}^*\|_2^2, \quad (25)$$

$$v := \text{Tr}(\text{Var}(\hat{\mathbf{x}}|\mathbf{x}^*)). \quad (26)$$

The motivation for analyzing this setting (even though in the proposed core procedure (3) & (21), we modify δ_t along the iterations and do not attempt to reach a minimizer per δ_t) is that insights on gains in bias and variance at the minimizer level hint that similar effects are possible even with limited optimizer iterations per δ_t . Also, in [38], it was shown that insights gained for Tikhonov prior generalize empirically to more complex priors. Specifically, they showed that BP typically leads to smaller bias but higher variance than LS. In what follows, the intuition that $0 < \delta < 1$ enables trading bias and variance is made formal.

Theorem 3.4. Consider the observation model (1) and estimating \mathbf{x}^* via minimization of (2) with $s(\mathbf{x}) = \frac{\beta}{2} \|\mathbf{D} \mathbf{x}\|_2^2$.

Assume that: (a) $\mathbf{A}^T \mathbf{A}$ and $\mathbf{D}^T \mathbf{D} \succ 0$ share eigenbasis; (b) the singular value of \mathbf{A} are in $(0, 1]$, and not all equal (common case); (c) $\eta = 0$ and $c = 1$. Then, $b_{BP} < b_{WLS} < b_{LS}$, and $v_{LS} < v_{WLS} < v_{BP}$.

Note that the margins in the theorem's inequalities depend on where δ is located in $(0, 1)$ (e.g., $v_{WLS} \rightarrow v_{LS}$ for $\delta \rightarrow 1$). This further motivates having a set $\{\delta_t\}$ instead of a fixed δ in (20) to gain flexibility in handling the error's bias/variance.

As for the convergence rates of first-order optimization algorithms (which affect the number of NFEs with practical DNN-based priors), recall that it can typically be characterized by the condition number, $\kappa(\cdot)$, of the problem's Hessian, which equals the ratio between the largest and smallest eigenvalues of the Hessian — the lower the better. Note, though, that for $m < n$, the Hessian of each of the three data-fidelity terms (LS, BP, WLS) is rank-deficient due to the ill-posedness of the observation model. Nevertheless, we have the following property on the component of their Hessian in the row-range of \mathbf{A} (a subspace in \mathbb{R}^n).

Claim 3.5. Assume that $\text{rank}(\mathbf{A}) = m$, the singular values of \mathbf{A} are not all equal, $\eta = 0$, and denote by $\mathbf{V} \in \mathbb{R}^{n \times m}$ an orthonormal basis for the row-range of \mathbf{A} . We have that

$$\kappa(\mathbf{V}^T \nabla_{\mathbf{x}}^2 \ell_{BP} \mathbf{V}) < \kappa(\mathbf{V}^T \nabla_{\mathbf{x}}^2 \ell_{WLS} \mathbf{V}) < \kappa(\mathbf{V}^T \nabla_{\mathbf{x}}^2 \ell_{LS} \mathbf{V}).$$

Namely, the flexibility in δ allows accelerating the reconstruction procedure compared to using only a LS based guidance. Leveraging this result to establish ranking claims for specific algorithms may follow the approaches of [39] (based on constraints on the null space of \mathbf{A} that are implicitly imposed by the signal prior term).

Algorithm 1 Denoising Diffusion with iterative Preconditioned Guidance (DDPG)

Require: $\epsilon_\theta(\cdot, t)$ (noise estimator), T , \mathbf{y} , \mathbf{A} , $\{\bar{\alpha}_t\}$, $\{\mu_t\}$, c , η , γ , ζ

- 1: **if** $\sigma_e > 0$ **then**
- 2: $\delta_t = \bar{\alpha}_t^\gamma$ and $w_t = \delta_t$
- 3: **else**
- 4: $\delta_t = 0$ and $w_t = 1$
- 5: **end if**
- 6: Initialize $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$
- 7: **for** t from T to 1 **do**
- 8: $\mathbf{x}_{0|t} = \frac{1}{\sqrt{\bar{\alpha}_t}} (\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t))$
- 9: $\mathbf{g}_{BP} = \mathbf{A}^T (\mathbf{A} \mathbf{A}^T + \eta \mathbf{I}_m)^{-1} (\mathbf{A} \mathbf{x}_{0|t} - \mathbf{y})$
- 10: $\mathbf{g}_{LS} = c \mathbf{A}^T (\mathbf{A} \mathbf{x}_{0|t} - \mathbf{y})$
- 11: $\tilde{\mathbf{x}}_{t-1} = \mathbf{x}_{0|t} - \mu_t ((1 - \delta_t) \mathbf{g}_{BP} + \delta_t \mathbf{g}_{LS})$
- 12: $\hat{\epsilon}_t = \frac{1}{\sqrt{1 - \bar{\alpha}_t}} (\mathbf{x}_t - \sqrt{\bar{\alpha}_t} \tilde{\mathbf{x}}_{t-1})$
- 13: $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$
- 14: $\mathbf{x}_{t-1} =$
- 15: $\sqrt{\bar{\alpha}_{t-1}} \tilde{\mathbf{x}}_{t-1} + \sqrt{1 - \bar{\alpha}_{t-1}} (w_t \sqrt{1 - \zeta} \hat{\epsilon}_t + \sqrt{\zeta} \epsilon_t)$
- 16: **end for**
- 17: **return** \mathbf{x}_0

3.3. Sampling scheme

As previously mentioned, in this subsection we turn IDPG (alternating (3) and (21)) into a DDM-based sampling scheme that empirically leads to better recoveries, especially in terms of perceptual quality. Injecting noise to the estimate in a certain iteration leads to an image that better matches the data on which the denoiser has been trained. Intuitively, this should improve the denoiser’s performance and also allow it to mitigate error propagation along the iterations, by masking artifacts with noise that will be removed. Accordingly, it allows using a smaller value of the regularization parameter η when handling noisy \mathbf{y} , which yields results with sharper details.

The sampling scheme that we propose is based on modification of the DDIM guided scheme, stated in (17). Namely, $\mathbf{x}_{0|t}$ is computed using a DDM’s noise estimator as in (16) and the novel guidance \mathbf{g}_{δ_t} is being used in lieu of the abstract gradient $\nabla_{\mathbf{x}} \ell$. However, we propose several important modifications of the additive noise terms $\hat{\sigma}_t \epsilon_\theta(\mathbf{x}_t; t) + \tilde{\sigma}_t \epsilon_t$ that appear in (17).

The first modification is inspired by [45]. Note that, contrary to the unconditional image generation task, the data-fidelity guidance can significantly modify the estimated signal $\tilde{\mathbf{x}}_{t-1} := \mathbf{x}_{0|t} - \mu_t \mathbf{g}_{\delta_t}$ compared to $\mathbf{x}_{0|t}$. Therefore, the *effective* predicted noise can be significantly different than $\epsilon_\theta(\mathbf{x}_t; t)$. Accordingly, based on the relation between $\epsilon_\theta(\mathbf{x}_t; t)$ and $\mathbf{x}_{0|t}$ in (16), the effective predicted noise takes

the form of

$$\hat{\epsilon}_t = \frac{1}{\sqrt{1 - \bar{\alpha}_t}} (\mathbf{x}_t - \sqrt{\bar{\alpha}_t} \tilde{\mathbf{x}}_{t-1}), \quad (27)$$

and we use it instead of $\epsilon_\theta(\mathbf{x}_t; t)$ for generating \mathbf{x}_{t-1} .

Similarly to this prior work, we also simplify the way of trading between the (effective) estimated noise and the stochastic noise. Instead of time-depended noise level $\tilde{\sigma}_t$ which affects $\hat{\sigma}_t = \sqrt{1 - \bar{\alpha}_{t-1} - \tilde{\sigma}_t^2}$ in a complicated way, we balance between the levels of $\hat{\epsilon}_t$ and ϵ_t using the simple to tune weight $\sqrt{1 - \zeta}$ and $\sqrt{\zeta}$, respectively, with a single hyperparameter $\zeta \in [0, 1]$.

Our key novel modification is the following. In the case of observation noise (i.e., $\sigma_e > 0$), the effective predicted noise $\hat{\epsilon}_t$ may not be useful in iterations where the overall guidance \mathbf{g}_{δ_t} is dominated by \mathbf{g}_{BP} , due to the sensitivity to noise of the BP guidance. Therefore, we scale it by δ_t — the weight that is given to \mathbf{g}_{LS} , which is more robust to observation noise.

We remark that our sampling scheme transfers smoothly to the noiseless case, where one can simply use pure BP guidance – or equivalently $\delta_t = 0$: it is only required to omit the last modification (i.e., avoid attenuating the noise injection by δ_t). As will be shown empirically, this scheme performs on par with BP-based DDM methods that are specifically devised for noiseless observations, and outperforms other methods.

The resulted DDM-based sampling algorithm, that we name Denoising Diffusion with iterative Preconditioned Guidance (DDPG), is summarized in Algorithm 1. The algorithm includes a simple way to set $\{\delta_t\}$ that is explained in the next subsection.

3.4. Implementation details

In the experimental section, we wish to test the approach (both DDPG and its core version IDPG) with powerful DDM’s denoisers trained and used in [7, 20]. These models are based on DDPM formulation [13, 29], that we use to reduce the hyperparameter tuning effort in all our examined tasks, and in particular setting $\{\delta_t\}$ for noisy observations. (Recall that in the noiseless case, we simply set $\delta_t = 0$).

As presented in Section 2, in DDPM there are the parameters $\{\bar{\alpha}_t\}$, which are determined by the hyperparameters $\{\beta_t\}$, and determine many other parameters (e.g., the model’s noise levels $\{\sigma_t\}$). These $\{\bar{\alpha}_t\}$ obey $0_+ \approx \bar{\alpha}_T < \dots < \bar{\alpha}_0 = 1$. We do not change $\{\beta_t\}$ and $\{\bar{\alpha}_t\}$ at all compared to previous methods that we compare with [16, 41]. In fact, we use $\{\bar{\alpha}_t\}$ to set $\delta_t = \bar{\alpha}_t^\gamma$ with a single positive scalar hyperparameter γ . This ensures that as t decreases, $\{\delta_t\}$ is monotonically increasing from 0 to 1 in rate that can be controlled by γ . Observe that smaller γ yields higher intermediate $\bar{\alpha}_t^\gamma$ and thus a larger portion of $\{\mathbf{W}_t\}$ closer to

identity matrix — or, equivalently, “faster” progress from BP to LS. This is beneficial for larger observation noise σ_e .

To conclude, when using $c = 1$ and $\mu_t = 1$, applying our approach requires tuning only η , γ and ζ (for DDPG). Note that, potentially, tuning more hyperparameters for our method can further boost its results.

4. Experiments

In this section, we examine the performance of our approach. We start with examining our core guidance, as reflected in IDPG, compared to its extreme cases: deterministic methods with pure BP and LS based guidance. Then, we compare both IDPG and DDPG to existing methods: DDRM [16], DPS [6], DDNM [41], and DiffPIR [45]. We consider the CelebA-HQ 1K and ImageNet 1K test sets, where for fair comparison all the methods use the same pre-trained DDM’s denoisers: the model from [20] trained on CelebA-HQ and the (unconditional) model from [7] trained on ImageNet. We use $T = 100$ iterations for each of the methods as in [41]. An exception is DPS, which requires $T = 1000$ iterations with much larger per-iteration complexity, making it extremely slow compared to others.¹ All the methods are initialized with the same random \mathbf{x}_T .

We consider image deblurring and super-resolution tasks that have been widely examined in the previous works, and for each method we use the hyperparameter settings suggested by its authors. The settings for our methods are stated in the supplementary material. Specifically, we consider super-resolution with bicubic downsampling and scale factor 4 and deblurring with Gaussian kernel as, e.g., in [16, 41]. However, we also examine challenging noisy deblurring with random motion kernels, as in [45]. This is an example for a common task that cannot be handled by DDRM [16], which is limited to separable kernels for which the SVD of \mathbf{A} can be efficiently computed and stored. This task cannot be addressed also by DDNM [41], which fails to tackle noisy observations.² We examine the noisy cases also in the two aforementioned tasks.

In the supplementary material, we demonstrate the advantages of our approach over recent task-specific DNNs, as well as its superior performance also with lower noise level than demonstrated here. In addition, we report there the impressive performance of our approach for sparse-view computed tomography (another task where computing the SVD of \mathbf{A} is infeasible).

¹Using the same hardware, DPS takes a dozen minutes per image compared to a couples of seconds of the other methods, and limiting it to 100 iterations fails to produce meaningful recoveries.

²Note that DDNM is tailored for $\sigma_e = 0$. We remark that DDNM+ that was proposed in [41] for handling noisy \mathbf{y} (via SVD) seems to be heavily tied to a specific downsampling task (without bicubic kernel) and does not support the considered tasks, as shown in the supplementary material. Our efforts to adapt it failed.

Table 1. Data-fidelity guidance ablation: super-resolution and deblurring PSNR [dB] (\uparrow) and LPIPS (\downarrow) results on CelebA-HQ 1K.

Task	PGM-LS	IDBP	IDPG (ours)
Bicub. SRx4 $\sigma_e=0$	32.40 / 0.127	32.66 / 0.111	32.66 / 0.111
Bicub. SRx4 $\sigma_e=0.05$	28.91 / 0.209	28.95 / 0.229	29.89 / 0.155
Gauss. Deb. $\sigma_e=0$	32.25 / 0.141	45.58 / 0.002	45.58 / 0.002
Gauss. Deb. $\sigma_e=0.05$	28.61 / 0.191	30.55 / 0.150	31.08 / 0.150
Gauss. Deb. $\sigma_e=0.1$	27.89 / 0.239	28.29 / 0.226	29.28 / 0.146



Figure 2. SRx4 with noise level 0.05. Top row, from left to right: original, upsampled observation, and our DDPG. Bottom row, from left to right: PGM-LS, IDBP, and our IDPG (baseline for DDPG).

4.1. Examining the core approach

As the main contribution of our paper is the novel guidance approach that smoothly traverses from BP to LS guidance, we start with examining our baseline IDPG method (alternating (3) and (21)) with IDBP [36] ($\delta_t = 0$) and a method with LS based updates ($\delta_t = 1$), denoted here by PGM-LS (a common PnP scheme [14]). This comparison essentially provides an ablation study for our core approach.

The results for bicubic super-resolution and Gaussian deblurring with various and without noise for the CelebA-HQ 1K test set are presented in Table 1. They show that IDPG consistently outperforms IDBP and PGM-LS. Figure 2 displays qualitative results, showing that IDPG has less artifacts than IDBP and finer details than PGM-LS. Yet, while IDPG provides high PSNR values its perceptual quality can be improved by using the sampling-based DDPG, as shown in this figure.

4.2. Comparison with other methods

We turn to compare the proposed approach to DDRM, DPS, DDNM and DiffPIR. We consider all deblurring and super-resolution tasks that have been described above with various noise levels. The results for the CelebA-HQ 1K test set and the ImageNet 1K test set are presented in Table 2 and Table 3, respectively. We present qualitative results for motion deblurring, Gaussian deblurring and super-resolution in Figures 1, 3 and 4, respectively. More results appear in the supplementary material.

Table 2. Super-resolution and deblurring PSNR [dB] (\uparrow) and LPIPS (\downarrow) results on CelebA-HQ 1K. N/A marks applicability limitation of: (1) DDNM to noiseless settings and (2) DDRM to settings where the SVD is given and stored. (More details in the text).

Task \ Method	DDRM	DPS (1000 NFEs)	DiffPIR	DDNM	IDPG (ours)	DDPG (ours)
Bicub. SRx4 $\sigma_e=0$	31.64 / 0.054	29.39 / 0.065	30.26 / 0.051	31.64 / 0.048	32.66 / 0.111	31.60 / 0.052
Bicub. SRx4 $\sigma_e=0.05$	29.26 / 0.090	27.49 / 0.086	27.44 / 0.085	N/A	29.89 / 0.155	29.39 / 0.105
Gauss. Deb. $\sigma_e=0$	42.49 / 0.006	31.25 / 0.055	32.97 / 0.041	45.56 / 0.002	45.58 / 0.002	45.46 / 0.002
Gauss. Deb. $\sigma_e=0.05$	30.53 / 0.074	27.75 / 0.084	28.89 / 0.074	N/A	31.08 / 0.150	30.41 / 0.068
Gauss. Deb. $\sigma_e=0.1$	28.79 / 0.088	26.67 / 0.097	27.59 / 0.083	N/A	29.28 / 0.146	29.18 / 0.080
Motion Deb. $\sigma_e=0.05$	N/A	19.63 / 0.227	27.96 / 0.102	N/A	29.73 / 0.134	29.02 / 0.082
Motion Deb. $\sigma_e=0.1$	N/A	19.64 / 0.231	26.23 / 0.132	N/A	27.86 / 0.166	27.74 / 0.099

Table 3. Super-resolution and deblurring PSNR [dB] (\uparrow) and LPIPS (\downarrow) results on Imagenet 1K. N/A marks applicability limitation of: (1) DDNM to noiseless settings and (2) DDRM to settings where the SVD is given and stored. (More details in the text).

Task \ Method	DDRM	DPS (1000 NFEs)	DiffPIR	DDNM	IDPG (ours)	DDPG (ours)
Bicub. SRx4 $\sigma_e=0$	27.38 / 0.270	25.56 / 0.236	26.99 / 0.225	27.45 / 0.245	27.20 / 0.326	27.41 / 0.255
Bicub. SRx4 $\sigma_e=0.05$	25.54 / 0.333	24.05 / 0.271	24.65 / 0.318	N/A	25.51 / 0.411	25.55 / 0.354
Gauss. Deb. $\sigma_e=0$	40.53 / 0.008	25.54 / 0.259	30.54 / 0.082	43.64 / 0.003	44.02 / 0.002	44.21 / 0.002
Gauss. Deb. $\sigma_e=0.05$	27.71 / 0.243	23.59 / 0.294	26.64 / 0.240	N/A	27.47 / 0.313	27.73 / 0.205
Motion Deb. $\sigma_e=0.05$	N/A	17.52 / 0.468	25.34 / 0.284	N/A	26.02 / 0.354	25.94 / 0.249

Examining the results, observe that, indeed, the sampling-based version of our approach (DDPG) consistently leads to better perceptual quality, both visually and as measured by LPIPS. Interestingly, while there is an inherent tradeoff between accuracy and perceptual quality [3], the PSNR reduction in DDPG compared to IDPG is rather moderate (and in some case DDPG even has slightly better PSNR than IDPG). One can observe that our approach typically outperforms other approaches in deblurring and is competitive in super-resolution.

Importantly, note that the only reference methods that are as flexible as our approach to the observation model are DPS, which is very slow, and DiffPIR. However, their PSNR is significantly lower than the PSNR of our DDPG. Recall that in many critical applications (e.g., medical imaging), accuracy, as measured by the classical PSNR metric, is extremely important.

5. Conclusion

In this paper, we presented a framework for solving linear inverse problems with DNN denoisers/diffusers, which uses a novel preconditioned data-fidelity guidance approach, based on traversing from back-projection steps to least squares steps, exploiting the advantages of each. We used the new approach with a computationally convenient trajectory within a “plug-and-play” optimization scheme and its DDM sampling-based counterpart, which we devised. The performance advantages of the new technique were shown in various deblurring and super-resolution settings, with and without observation noise. (See the supplementary material for computed tomography experiments as well). As a direction for future research, one can try to learn the preconditioners $\{\mathbf{W}_t\}$ instead of designing them.




Figure 3. Gaussian deblurring with noise level 0.05. Top row, from left to right: original, observation, and our IDPG (baseline for DDPG). Bottom row, from left to right: DPS, DiffPIR, and our DDPG.




Figure 4. SRx4 with noise level 0.05. Top row, from left to right: original, upsampled observation, and our IDPG (baseline for DDPG). Bottom row, from left to right: DDRM, DiffPIR, and our DDPG.

tioners $\{\mathbf{W}_t\}$ instead of designing them.

Acknowledgment. The work is partially supported by ISF grant no. 1940/23.

References

- [1] Shady Abu Hussein, Tom Tirer, and Raja Giryes. Correction filter for single image super-resolution: Robustifying off-the-shelf deep super-resolvers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1428–1437, 2020. **1**
- [2] Shady Abu-Hussein, Tom Tirer, and Raja Giryes. Adir: Adaptive diffusion for image reconstruction. *arXiv preprint arXiv:2212.03221*, 2022. **2**
- [3] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6228–6237, 2018. **8**
- [4] Stanley H Chan, Xiran Wang, and Omar A Elgendy. Plug-and-play admm for image restoration: Fixed-point convergence and applications. *IEEE Transactions on Computational Imaging*, 3(1):84–98, 2016. **2**
- [5] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*, 35:25683–25696, 2022. **2, 16, 17**
- [6] Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations*, 2023. **1, 2, 3, 7, 15, 16**
- [7] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. **2, 6, 7**
- [8] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015. **1**
- [9] Bradley Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011. **3**
- [10] Michael Elad, Bahjat Kawar, and Gregory Vaksman. Image denoising: The deep learning revolution and beyond—a survey paper. *SIAM Journal on Imaging Sciences*, 16(3):1594–1654, 2023. **3**
- [11] Martin Genzel, Ingo Gühring, Jan Macdonald, and Maximilian März. Near-exact recovery for tomographic inverse problems via deep learning. In *International Conference on Machine Learning*, pages 7368–7381. PMLR, 2022. **2**
- [12] Magnus Rudolph Hestenes and Eduard Stiefel. *Methods of conjugate gradients for solving linear systems*. 1952. **3**
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. **2, 3, 6**
- [14] Ulugbek S Kamilov, Charles A Bouman, Gregory T Buzzard, and Brendt Wohlberg. Plug-and-play methods for integrating physical and learned models in computational imaging: Theory, algorithms, and applications. *IEEE Signal Processing Magazine*, 40(1):85–97, 2023. **2, 7**
- [15] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35:26565–26577, 2022. **3**
- [16] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022. **2, 3, 4, 6, 7, 15**
- [17] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1833–1844, 2021. **16**
- [18] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017. **1**
- [19] Gongye Liu, Haoze Sun, Jiayi Li, Fei Yin, and Yujiu Yang. Accelerating diffusion models for inverse problems through shortcut sampling. *arXiv preprint arXiv:2305.16965*, 2023. **2**
- [20] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022. **2, 6, 7**
- [21] Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving inverse problems with diffusion models. *arXiv preprint arXiv:2305.04391*, 2023. **2**
- [22] Jean-Jacques Moreau. Proximité et dualité dans un espace hilbertien. *Bull. Soc. Math. France*, 93(2):273–299, 1965. **2**
- [23] Jorge Nocedal and Stephen J Wright. *Numerical optimization*. Springer, 1999. **4**
- [24] Edward T Reehorst and Philip Schniter. Regularization by denoising: Clarifications and new interpretations. *IEEE transactions on computational imaging*, 5(1):52–67, 2018. **2**
- [25] Yaniv Romano, Michael Elad, and Peyman Milanfar. The little engine that could: Regularization by denoising (RED). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844, 2017. **2**
- [26] Anand P Sabulal and Srikrishna Bhashyam. Joint sparse recovery using deep unfolding with application to massive random access. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5050–5054. IEEE, 2020. **2**
- [27] Marion Savanier, Emilie Chouzenoux, Jean-Christophe Pesquet, and Cyril Riddell. Deep unfolding of the dbfb algorithm with application to roi ct imaging with limited angular density. *IEEE Transactions on Computational Imaging*, 2023. **2**
- [28] Assaf Shocher, Nadav Cohen, and Michal Irani. “zero-shot” super-resolution using deep internal learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3118–3126, 2018. **1**
- [29] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. **3, 6**
- [30] Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse

- problems. In *International Conference on Learning Representations*, 2023. 2, 3, 4
- [31] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019. 2, 3
- [32] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. *arXiv preprint arXiv:2111.08005*, 2021. 2
- [33] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. 3, 16
- [34] Jian Sun, Wenfei Cao, Zongben Xu, and Jean Ponce. Learning a convolutional neural network for non-uniform motion blur removal. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 769–777, 2015. 1
- [35] Yu Sun, Brendt Wohlberg, and Ulugbek S Kamilov. An online plug-and-play algorithm for regularized image reconstruction. *IEEE Transactions on Computational Imaging*, 5(3):395–408, 2019. 2
- [36] Tom Tirer and Raja Giryes. Image restoration by iterative denoising and backward projections. *IEEE Transactions on Image Processing*, 28(3):1220–1234, 2018. 2, 7
- [37] Tom Tirer and Raja Giryes. Super-resolution via image-adapted denoising CNNs: Incorporating external and internal learning. *IEEE Signal Processing Letters*, 26(7):1080–1084, 2019. 1, 2
- [38] Tom Tirer and Raja Giryes. Back-projection based fidelity term for ill-posed linear inverse problems. *IEEE Transactions on Image Processing*, 29(1):6164–6179, 2020. 2, 3, 5, 12
- [39] Tom Tirer and Raja Giryes. On the convergence rate of projected gradient descent for a back-projection based objective. *SIAM Journal on Imaging Sciences*, 14(4):1504–1531, 2021. 2, 3, 5
- [40] Singanallur V Venkatakrishnan, Charles A Bouman, and Brendt Wohlberg. Plug-and-play priors for model based reconstruction. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 945–948. IEEE, 2013. 2
- [41] Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. In *International Conference on Learning Representations*, 2023. 2, 3, 4, 6, 7, 15, 16
- [42] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5728–5739, 2022. 16
- [43] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017. 1
- [44] Kai Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep cnn denoiser prior for image restoration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3929–3938, 2017. 2
- [45] Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jie Zhang Cao, Bihan Wen, Radu Timofte, and Luc Van Gool. Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1219–1229, 2023. 1, 2, 3, 4, 6, 7, 15, 16
- [46] Jenny Zukerman, Tom Tirer, and Raja Giryes. Bp-dip: A backprojection based deep image prior. In *2020 28th European Signal Processing Conference (EUSIPCO)*, pages 675–679, 2021. 2