

Integrating Reweighted Least Squares with Plug-and-Play Diffusion Priors for Noisy Image Restoration

Ji Li¹, Chao Wang^{2*}

¹Academy for Multidisciplinary Studies, Capital Normal University, Beijing, China

²College of Science, China Agricultural University, Beijing, China
matliji@163.com, wywwwnx@163.com

Abstract

Existing plug-and-play image restoration methods typically employ off-the-shelf Gaussian denoisers as proximal operators within classical optimization frameworks based on variable splitting. Recently, denoisers induced by generative priors have been successfully integrated into regularized optimization methods for image restoration under Gaussian noise. However, their application to non-Gaussian noise—such as impulse noise—remains largely unexplored. In this paper, we propose a plug-and-play image restoration framework based on generative diffusion priors for robust removal of general noise types, including impulse noise. Within the maximum a posteriori (MAP) estimation framework, the data fidelity term is adapted to the specific noise model. Departing from the conventional least-squares loss used for Gaussian noise, we introduce a generalized Gaussian scale mixture-based loss, which approximates a wide range of noise distributions and leads to an ℓ_q -norm fidelity term. This optimization problem is addressed using an iteratively reweighted least squares (IRLS) approach, wherein the proximal step involving the generative prior is efficiently performed via a diffusion-based denoiser. Experimental results on benchmark datasets demonstrate that the proposed method effectively removes non-Gaussian impulse noise and achieves superior restoration performance.

Code — <https://github.com/Chilie/IRLS-PnPDP>

Extended version — <https://arxiv.org/abs/2511.06823>

Introduction

Image restoration, encompassing tasks such as denoising, deblurring, super-resolution and inpainting, is an important and longstanding problem in computer vision and signal processing. It is typically formulated as an inverse problem, aiming to recover the latent clean image \mathbf{x} from a corrupted observation \mathbf{y} , modeled by the following system:

$$\mathbf{y} = \mathcal{D}(\mathbf{Ax}), \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^n$ denotes the clean image, $\mathbf{A} \in \mathbb{R}^{m \times n}$ models the degradation operator, $\mathcal{D} : \mathbb{R}^m \rightarrow \mathbb{R}^m$ represents

the noise corruption process. The system (1) is generally ill-posed, as the dimension of measurement m is often less than that of the latent image n .

To stabilize the recovery of \mathbf{x} , variational optimization approaches introduce regularization that incorporates prior knowledge about the solution. A common formulation is:

$$\min_{\mathbf{x}} \mathcal{S}(\mathbf{Ax}, \mathbf{y}) + \lambda \mathcal{R}(\mathbf{x}), \quad (2)$$

where $\mathcal{S}(\mathbf{Ax}, \mathbf{y})$ measures the fidelity to the observed data, $\mathcal{R}(\mathbf{x})$ enforces the prior information about the statistics of the latent image (Chambolle et al. 2010; Beck and Teboulle 2009; Chan et al. 2006; Rudin, Osher, and Fatemi 1992; Figueiredo and Nowak 2001). Notably, this formulation corresponds to a maximum a posteriori (MAP) estimation of the posterior density $p(\mathbf{x}|\mathbf{y})$, where $\mathcal{S}(\mathbf{Ax}, \mathbf{y}) = -\log p(\mathbf{y}|\mathbf{x})$ and $\lambda \mathcal{R}(\mathbf{x}) = -\log p(\mathbf{x})$ respectively.

The fidelity term $\mathcal{S}(\mathbf{Ax}, \mathbf{y})$ should ideally reflect noise distribution, while regularization term $\mathcal{R}(\mathbf{x})$ should approximate the logarithm of image prior. The plug-and-play (PnP) framework and regularization-by-denoising (RED) (Reehorst and Schniter 2018; Romano, Elad, and Milanfar 2017) have emerged as widely adopted solutions for various image restoration tasks. Within the framework, regularization is implicitly enforced via a denoiser, which can be realized using classical methods such as BM3D (Dabov et al. 2007), deep neural network-based (Zhang et al. 2017), and generative models-based denoisers (Chung et al. 2022; Zhu et al. 2023; Martin et al. 2024).

Given a noiseless measurement $\mathbf{y}_{\text{nl}} := \mathbf{Ax}$, the noise corruption process can be expressed as $\mathbf{y} = \mathbf{y}_{\text{nl}} + \mathbf{s}$. When \mathbf{s} is a pixel-independent Gaussian, the likelihood $p(\mathbf{y}|\mathbf{x})$ leads to a least-squares loss in (2). For non-Gaussian, pixel-dependent noise, the exact form of $p(\mathbf{y}|\mathbf{x})$ may be unknown. To address this, we approximate the noise distribution using a generalized Gaussian scale mixture (gGSM) model:

$$P(s; \delta, q) = \frac{q}{2\delta\Gamma(1/q)} \exp\left(-\frac{|s|^q}{\delta^q}\right), \quad (3)$$

where δ, q are the distribution parameters, and $\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt$ is the Gamma function. The gGSM contains two special known cases: $q = 1$ corresponds to the Laplacian scale mixture model

$$P(s; \theta) = \frac{1}{2\theta} \exp\left(-\frac{|s|}{\theta}\right) \quad (4)$$

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

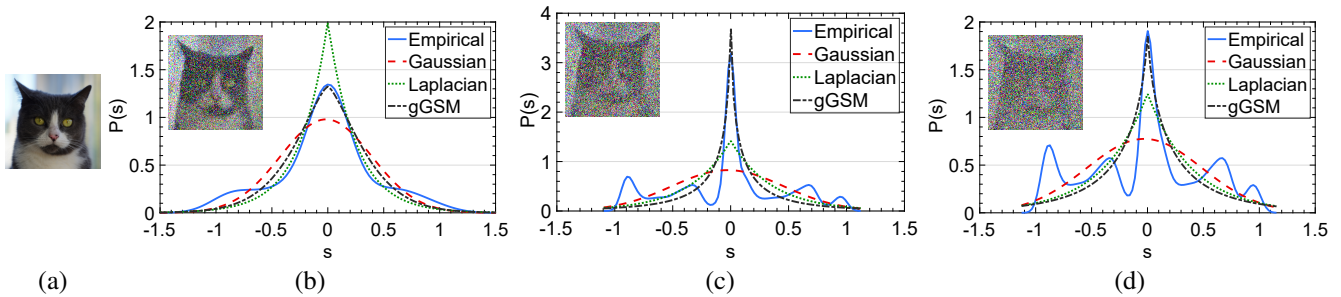


Figure 1: Empirical distribution of impulse noise and its fitted approximation densities (including Laplacian, Gaussian and gGSM) using maximum likelihood estimation. (a) The clean image; (b) Noise level $r_{sp} = 0.5$, the fitted parameter $\theta = 0.25$, $\delta = 0.44$, $q = 1.5$. (c) Noise level $r_{sp} = 0.7$, the fitted parameter $\theta = 0.35$, $\delta = 0.05$, $q = 0.5$. (d) Noise level $r_{sp} = 0.8$, the fitted parameter $\theta = 0.40$, $\delta = 0.20$, $q = 0.7$. For (b)-(d), the noisy image displays on the top left corner.

with parameter θ and $q = 2$ yields the Gaussian density.

As a motivating example, we consider impulse noise, especially the salt-and-pepper noise (Windyga 2001), which frequently arises in data acquisition and transmission due to faulty sensors or sharp and sudden disturbances in analog-to-digital conversion. Salt-and-pepper noise sets random pixels to the minimum or maximum intensity with equal probability, controlled by a noise level r_{sp} . Figure 1 compares the empirical distribution of impulse noise with several approximation models, including Gaussian, Laplacian, and gGSM. Among them, gGSM provides the best fit for modeling the noise distribution, typically with $q \in (0, 1]$.

Given that the gGSM model leads to an ℓ_q -norm ($0 < q \leq 2$) data fidelity term in (2), we aim to solve the corresponding minimization problem using a plug-and-play approach with generative priors. Specifically, we propose an efficient iterative reweighted least squares (IRLS) algorithm to solve a surrogate problem derived from the ℓ_q -norm formulation. This surrogate is then addressed using a forward-backward splitting method, where the proximal step is performed via a denoiser induced by a generative diffusion model.

To summarize, our main contributions are as follows:

1. We introduce a plug-and-play image restoration framework for general noise removal, formulating an ℓ_q -norm minimization problem using the gGSM model to characterize noise distributions;
2. We propose an IRLS-based majorization-minimization (MM) scheme (Sun, Babu, and Palomar 2016) to solve the ℓ_q -norm sub-problem. This is efficiently addressed using forward-backward splitting, with the proximal step handled via a generative diffusion denoiser;
3. To mitigate potential division-by-zero issue in IRLS, we propose a minor modification over IRLS with a decreasing hyperparameter, following (Chartrand and Yin 2008).

Related Works

Traditional Method for Image Restoration

Traditional image restoration methods are typically distinguished by their choice of data fidelity term and image prior. A widely used class of priors involves sparsity-promoting

regularization, such as total variation (TV) and its variants (Blomgren et al. 1997; Vogel and Oman 1996), or sparsity in wavelet-transformed domains (Figueiredo and Nowak 2001). These priors promote smoothness in the restored image, prioritizing the recovery of low-frequency components. However, due to a mismatch between hand-crafted regularizations and the true image prior, the recovered images are generally suboptimal. The choice of data fidelity term is often noise-model-dependent. The ℓ_2 -norm is commonly used for Gaussian noise, while the ℓ_1 -norm is more suitable for sparse noise. For impulse noise, the ℓ_0 -norm is popular but challenging to optimize; it is often replaced by continuous concave approximations (Koshelev and Lefkimmiatis 2023; Zhang et al. 2015; Ochs et al. 2015).

Plug-and-play Methods

Since introduced in (Venkatakrishnan, Bouman, and Wohlberg 2013), the idea of plug-and-play has gained significant attention due to its flexibility and effectiveness in addressing diverse image restoration problems. With fixed data fidelity term, the regularized optimization is generally solved by a splitting method. The regularization term is handled implicitly through a denoising operation, which acts as a proximal step. There are many effective denoiser candidates can be used, such as BM3D (Dabov et al. 2007) denoiser and DnCNN/DRUNet denoiser (Zhang et al. 2017, 2021). The incorporation of the denoiser into regularized optimization has been adopted to solve image restoration (Zhang et al. 2020; Ryu et al. 2019; Venkatakrishnan, Bouman, and Wohlberg 2013; Kamilov, Mansour, and Wohlberg 2017), and medical image reconstruction (Eksioğlu 2016; Metzler, Maleki, and Baraniuk 2015, 2016). The RED is another way to explicitly formulate the regularization (Romano, Elad, and Milanfar 2017).

Plug-and-play with Generative Prior

Recently, powerful diffusion generative models have been leveraged as priors for image restoration tasks. Rather than modeling image density directly, diffusion models learn to approximate the score function $\nabla_x \log p(\mathbf{x})$ (Song and Ermon 2019; Song et al. 2020). Among generative models, diffusion models have demonstrated superior performance

compared to alternatives such as variational autoencoders (VAEs) and generative adversarial networks (GANs), making them highly appealing for plug-and-play image restoration. In zero-shot learning settings, off-the-shelf diffusion models can be used for observation-guided conditional sampling from the posterior $p(\mathbf{x}|\mathbf{y})$. A key step involves approximating the conditional score $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t|\mathbf{y})$:

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

Thus, estimating $\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t)$ becomes critical. Chung et al. (2022) leveraged estimate $p(\mathbf{y}|\mathbf{x}_t) \simeq p(\mathbf{y}|\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t])$, with $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$ estimated by Tweedie’s formula. This approach requires backpropagation through the diffusion model, which is computationally intensive. Moreover, their approximation has been considered coarse. Song et al. (2022) proposed IIGDM method with a Gaussian estimate of unknown $p(\mathbf{x}_0|\mathbf{x}_t)$. However, this method is limited to linear degradation models with Gaussian noise and does not extend naturally to non-Gaussian noise.

An alternative approach treats image restoration not as conditional generation but as optimization. DiffPIR (Zhu et al. 2023) applies half-quadratic splitting (HQS) to (2), using a diffusion model as a denoiser, thereby avoiding backpropagation through the model. However, their method requires multiple iterations to solve each subproblem. Li and Wang (2024) incorporated diffusion generation into RAAR to solve Fourier phase retrieval. Along this line, PnP-Flow (Martin et al. 2024) replaces HQS with forward-backward splitting (Parikh, Boyd et al. 2014), incorporating flow-based generative models instead of diffusion models. Additionally, Mardani et al. (2024) derived an explicit regularizer for diffusion priors, but the computation of its gradient is prohibitively expensive.

Supervised Learning for Image Restoration

The supervised learning solutions for image restoration typically train end-to-end convolutional neural networks using paired datasets (Gilton, Ongie, and Willett 2019; Meinhardt et al. 2017; Jin et al. 2017). For linear image restoration, the optimization unrolling scheme helps design an efficient network architecture, which enables better interpretation and training datasets reduction (Adler and Öktem 2018; Aggarwal, Mani, and Jacob 2018). These methods require retraining when task changes, hence lacking flexibility for solving a series of tasks, even the images share the same density.

Methodology

Before presenting our method, we give a brief introduction of the diffusion generative model, the formulated diffusion model regularized optimization and the majorization-minimization algorithm for ℓ_q -norm minimization.

Regularization via Diffusion Generative Model

The diffusion model comprises of two processes: forward process and generative process. The forward process progressively corrupts clean data \mathbf{x}_0 using a sequence of pre-configured noise scales. The seminar work (Song et al. 2020) provided a continuous modeling of the diffusion generative

model though the stochastic differential equation, we refer the reader to it for the details. For a self-contained organization, we follow the configurations of the denoising diffusion probability model (DDPM). Following (Ho, Jain, and Abbeel 2020), the process is formulated as a Markov chain: $p(\mathbf{x}_i|\mathbf{x}_{i-1}) = \mathcal{N}(\mathbf{x}_i; \sqrt{1-\beta_i}\mathbf{x}_{i-1}, \beta_i\mathbf{I})$, where $\{\beta_i\}_{i=1}^N$ is a sequence of noise levels with $0 < \beta_i < 1$. By recursion, one can show that $p(\mathbf{x}_i|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_i; \sqrt{\alpha_i}\mathbf{x}_0, (1-\alpha_i)\mathbf{I})$, where $\alpha_i := \prod_{j=1}^i (1-\beta_j)$.

To generate new samples, we aim to reverse this process. However, the exact reverse distribution $q(\mathbf{x}_{i-1}|\mathbf{x}_i)$ is intractable. Therefore, Ho, Jain, and Abbeel (2020) proposed to learn a Gaussian approximation $p_\theta(\mathbf{x}_{i-1}|\mathbf{x}_i)$ to approximate the true conditional probabilities $q(\mathbf{x}_{i-1}|\mathbf{x}_i)$. We know that

$$q(\mathbf{x}_{i-1}|\mathbf{x}_i, \mathbf{x}_0) = p(\mathbf{x}_i|\mathbf{x}_{i-1}, \mathbf{x}_0) \frac{p(\mathbf{x}_{i-1}|\mathbf{x}_0)}{p(\mathbf{x}_i|\mathbf{x}_0)} \text{ is a Gaussian.} \quad (5)$$

Specially, the mean and variance are

$$\begin{aligned} \tilde{\mu}(\mathbf{x}_i, \mathbf{x}_0) &= \frac{\sqrt{1-\beta_i}(1-\alpha_{i-1})}{1-\alpha_i} \mathbf{x}_i + \frac{\sqrt{\alpha_{i-1}}}{1-\alpha_i} \mathbf{x}_0, \\ \tilde{\sigma}^2 &= \frac{1-\alpha_{i-1}}{1-\alpha_i} \cdot \beta_i. \end{aligned}$$

Since \mathbf{x}_0 is unknown during sampling, we estimate \mathbf{x}_0 using \mathbf{x}_t . The estimate of $p(\mathbf{x}_0|\mathbf{x}_t)$ with given $p(\mathbf{x}_t|\mathbf{x}_0)$ is generally intractable, however, the mean of $p(\mathbf{x}_0|\mathbf{x}_i)$ can be computed using the Tweedie’s formula

$$\mathbb{E}[\mathbf{x}_0|\mathbf{x}_i] = \frac{\mathbf{x}_i + (1-\alpha_i)\nabla_{\mathbf{x}_i} \log p(\mathbf{x}_i)}{\sqrt{\alpha_i}}. \quad (6)$$

Diffusion model aims to train a network $s_\theta(\mathbf{x}_i, i)$ to predict the score function $\nabla_{\mathbf{x}_i} \log p(\mathbf{x}_i)$. Alternatively, DDPM trained a network $\epsilon_\theta(\mathbf{x}_i, i) \simeq -\sqrt{1-\alpha_i}\nabla_{\mathbf{x}_i} \log p(\mathbf{x}_i)$. Hence

$$\mathbb{E}[\mathbf{x}_0|\mathbf{x}_i] \simeq \frac{\mathbf{x}_i - \sqrt{1-\alpha_i}\epsilon_\theta(\mathbf{x}_i, i)}{\sqrt{\alpha_i}}. \quad (7)$$

Note that $\mathbf{x}_0 = \frac{\mathbf{x}_i - \sqrt{1-\alpha_i}\mathbf{z}}{\sqrt{\alpha_i}}$, comparing the formula to (7), hence the training loss

$$\mathcal{L} := \sum_i \frac{1-\alpha_i}{\alpha_i} \|\epsilon_\theta(\mathbf{x}_i, i) - \mathbf{z}\|_2^2, \quad (8)$$

where $\mathbf{x}_i = \sqrt{\alpha_i}\mathbf{x}_0 + \sqrt{1-\alpha_i}\mathbf{z}$. When deep diffusion model is available, we can perform DDPM, DDIM (Song, Meng, and Ermon 2021), and other method to perform the unconditional sampling.

The training loss (8) actually can be expressed as

$$\mathcal{L} := \sum_i \|\text{Denoise}[\mathbf{x}_i] - \mathbf{x}_0\|_2^2,$$

where $\text{Denoise}[\mathbf{x}_i]$ is the MMSE denoiser, which is exactly the equation (7). Using this in a regularized optimization framework, we propose to solve:

$$\min_{\mathbf{x}} \frac{1}{\lambda} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_q^q + \sum_i \|\text{Denoise}[\mathbf{x}_i] - \mathbf{x}\|_2^2, \quad (9)$$

where $\mathbf{x}_i = \sqrt{\alpha_i}\mathbf{x} + \sqrt{1-\alpha_i}\mathbf{z}$ and $\lambda = \delta^q$. The parameters δ and q are estimated from the gGSM model for the considered noise process of image restoration.

Majorization-minimization Algorithm

We aim to solve the problem with the following formulation:

$$\min \mathcal{Q}(\mathbf{x}) := \frac{1}{\lambda} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_q^q + \mathcal{R}(\mathbf{x}), \quad (10)$$

where the definition $\|\mathbf{v}\|_q^q = \sum_{i=1}^n |v_i|^q$ ($0 < q \leq 2$) for a vector $\mathbf{v} \in \mathbb{R}^n$. Direct minimization is difficult due to the non-smooth and non-convex nature of the ℓ_q -norm. To address this, we adopt a majorization-minimization (MM) approach (Sun, Babu, and Palomar 2016), which iteratively minimizes an upper-bound surrogate function. At each iteration k , we construct a surrogate $\mathcal{Q}^k(\mathbf{x}; \mathbf{x}_k)$ such that:

$$\mathcal{Q}(\mathbf{x}) \leq \mathcal{Q}^k(\mathbf{x}; \mathbf{x}_k), \quad \mathcal{Q}(\mathbf{x}_k) = \mathcal{Q}^k(\mathbf{x}_k; \mathbf{x}_k). \quad (11)$$

Minimizing \mathcal{Q}^k guarantees:

$$\mathcal{Q}(\mathbf{x}_{k+1}) \leq \mathcal{Q}^k(\mathbf{x}_{k+1}; \mathbf{x}_k) \leq \mathcal{Q}^k(\mathbf{x}_k; \mathbf{x}_k) = \mathcal{Q}(\mathbf{x}_k), \quad (12)$$

ensuring that $\mathcal{Q}(\mathbf{x}_k)$ is non-increasing. Under mild conditions, sequence \mathbf{x}_k converges to a stationary point of (10).

Then we majorize the $\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_q^q$ loss term.

Lemma 1. For any $x, y \in \mathbb{R}$ with $y \neq 0$ and $0 < q \leq 2$, the following inequality holds:

$$|x|^q \leq \frac{q}{2} |y|^{q-2} x^2 + \frac{2-q}{2} |y|^q. \quad (13)$$

Proof. Consider the function $f(x) = |x|^{q/2}$, by mean value theorem, the result holds by taking $x \leftarrow x^2$. \square

Lemma 2. For any $\mathbf{x} \in \mathbb{R}^d$, consider the function $f(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_q^q$, then we have

$$\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_q^q \leq \frac{q}{2} \|\mathbf{W}_k \circ (\mathbf{A}\mathbf{x} - \mathbf{y})\|_2^2 + \text{const}, \quad (14)$$

where the constant term is independent of \mathbf{x} . $\mathbf{W}_k := ((\mathbf{A}\mathbf{x}_k - \mathbf{y})^2 + \epsilon_k)^{\frac{q-2}{4}}$ is the reweighted matrix and \circ denotes the element-wise (Hadamard) product.

Iteratively Reweighted Least-squares

Using the MM approach, we solve the regularized least-squares minimization at each iteration:

$$\min_{\mathbf{x}} \frac{q}{2\lambda} \|\mathbf{W}_k \circ (\mathbf{A}\mathbf{x} - \mathbf{y})\|_2^2 + \sum_i \|\text{Denoise}[\mathbf{x}_i] - \mathbf{x}\|_2^2. \quad (15)$$

To solve (15), we adopt the forward-backward splitting method (Parikh, Boyd et al. 2014), which alternates between a gradient descent step and a proximal step:

$$\begin{aligned} \mathbf{x}_t &= \mathbf{x}_t - \frac{s_t \eta_k q}{\lambda} \mathbf{A}^T (\mathbf{W}_k^2 \circ (\mathbf{A}\mathbf{x}_t - \mathbf{y})) \\ \mathbf{x}_{t+1} &= \min_{\mathbf{x}} \frac{1}{2\eta_k} \|\mathbf{x} - \mathbf{x}_t\|_2^2 + \sum_i \|\text{Denoise}[\mathbf{x}_i] - \mathbf{x}\|_2^2, \end{aligned} \quad (16)$$

where s_t is a tunable stepsize. For convergence, it is typically required that $s_t \eta_k < \frac{2}{L}$, where L is the Lipschitz constant of the gradient of the least-squares term.

Solving the proximal step exactly requires computing the gradient of the regularization term, which involves back-propagation through the denoiser—this is computationally expensive. While some specific formulations allow for gradient simplification or avoidance (Mardani et al. 2024), evaluating the full summation remains costly. To mitigate this, we adopt a PnP strategy: approximate the proximal step using a denoiser derived from the diffusion model.

To align the noise level η_k with the noise distribution expected by the diffusion model, we set $\eta_k = (1 - \alpha_k)/\alpha_k$. However, the intermediate iterate \mathbf{x}_t generally does not match the input distribution for which the diffusion model ϵ_θ is trained. A common remedy is to renoise \mathbf{x}_t by adding Gaussian noise:

$$\begin{aligned} \mathbf{x}'_t &= \sqrt{\alpha_k} \mathbf{x}_t + \sqrt{1 - \alpha_k} \mathbf{z} \\ \mathbf{x}_{t+1} &= \text{Denoise}[\mathbf{x}'_t]. \end{aligned}$$

The ‘renoisng step’ is often referred to time reversal, has been widely used in conditional generation for image restoration (Choi et al. 2021; Wang, Yu, and Zhang 2022).

Algorithm 1: Image restoration with general noise removal via plug-and-play diffusion prior.

Require: Outer iterations T , pretrained diffusion model ϵ_θ and the noise schedule $\{\alpha_t\}$, hyperparameter q ($0 < q \leq 2$), pre-defined noise level $\{\eta_t = \frac{1-\alpha_t}{\alpha_t}\}$, perturbation sequences $\{\epsilon_t\}$, inter iteration T_{inter}

Ensure: Estimated image \mathbf{x}_0 .

- 1: Set $\mathbf{x}_T = \mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$
- 2: **for** $t = T : -1 : 1$ **do**
- 3: Set $\mathbf{x}_0 = \mathbf{x}_t$ and the reweighted matrix $\mathbf{W}_t = ((\mathbf{A}\mathbf{x}_t - \mathbf{y})^2 + \epsilon_t)^{q-2}$
- 4: $\% \%$ Solve MM (15) via forward-backward splitting
- 5: **for** $i = 0 : 1 : T_{inter} - 1$ **do**
- 6: $\% \%$ Gradient descent step:
 $\mathbf{x}_i = \mathbf{x}_i - s_i \eta_t \mathbf{A}^T (\mathbf{W}_t (\mathbf{A}\mathbf{x}_i - \mathbf{y}))$
- 7: $\% \%$ Denoising step:

$$\begin{aligned} \mathbf{x}'_i &= \sqrt{\alpha_t} \mathbf{x}_i + \sqrt{1 - \alpha_t} \mathbf{z} \\ \mathbf{x}_{i+1} &= \frac{\mathbf{x}'_i - \sqrt{1 - \alpha_t} \epsilon_\theta(\mathbf{x}'_i, t)}{\sqrt{\alpha_t}} \end{aligned}$$

- 8: **end for**
 - 9: $\% \%$ Set the update $\mathbf{x}_{t-1} = \mathbf{x}_{T_{inter}}$
 - 10: **end for**
-

Overflow avoidance via decreasing perturbation ϵ_k To prevent overflow, we introduce a small perturbation $\epsilon_k > 0$ in the weighted matrix $\mathbf{W}_k = ((\mathbf{A}\mathbf{x}_k - \mathbf{y})^2 + \epsilon_k)^{\frac{q-2}{2}}$. Moreover, the choice of ϵ_k affects both convergence speed and the sparsity of the solution (due to the nonconvexity of the ℓ_q term). Following (Chartrand and Yin 2008), we employ a decreasing perturbation schedule ϵ_k , which both stabilizes early iterations and improves sparsity in the later stage.

To summarize, we propose an image restoration method that addresses general noise via an ℓ_q -norm fidelity term, combined with a plug-and-play diffusion prior. The resulting




Figure 2: Visualization of different methods on two tasks with impulse noise for FFHQ. The DPS equipped with the IRLS loss produces artifact-corrupted restoration.

iteratively reweighted least-squares problem is solved using forward-backward splitting. The proximal step leverages the denoiser induced by a diffusion generative model, with careful renoising to ensure effective denoising. The complete algorithm is presented in Algorithm 1.

Updated Losses for DPS Method

The ℓ_q -norm data fidelity term can also be incorporated into the existing Diffusion Posterior Sampling (DPS) framework (Chung, Lee, and Ye 2023) to address image restoration with general noise corruption.

To adapt DPS for impulse noise, we consider two straightforward modifications to the standard DPS pipeline, which is originally designed for Gaussian noise. These modifications alter the measurement guidance step in DPS, where the transition from unconditional to conditional sampling occurs. The first modification directly applies gradient descent on the ℓ_q -norm fidelity term during the measurement guidance step. Specifically, we perform the following updates:

$$\begin{aligned} \mathbf{x}_{t-1} &= \text{UncondSample}(\mathbf{x}_t, \epsilon_\theta(\mathbf{x}_t, t)) \\ \mathbf{x}_{t-1} &= \mathbf{x}_{t-1} - \rho_t \nabla_{\mathbf{x}_t} \|\mathbf{A}\hat{\mathbf{x}}_0(\mathbf{x}_t) - \mathbf{y}\|_q^q, \end{aligned}$$

where $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ is the estimate of $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$ using Tweedie’s formula. However, this modification suffers from poor performance due to the nonconvexity of ℓ_q -norm (especially for $q < 1$), leading to unstable and suboptimal reconstructions.

Alternatively, following the IRLS of our approach, one can perform the measurement guidance step:

$$\mathbf{x}_{t-1} = \mathbf{x}_{t-1} - \rho_t \nabla_{\mathbf{x}_t} \|\mathbf{W}_t \circ (\mathbf{A}\hat{\mathbf{x}}_0(\mathbf{x}_t) - \mathbf{y})\|_2^2, \quad (17)$$

where $\mathbf{W}_t = ((\mathbf{A}\mathbf{x}_t - \mathbf{y})^2 + \epsilon_t)^{\frac{q-2}{4}}$. This modified DPS method performs better than the naive ℓ_q gradient approach and yields feasible restoration results. However, it still suffers from visible artifacts and reduced fidelity compared to our full IRLS-based plug-and-play method. Comparative results are provided in Figure 2.

Applicability for Other Generative Models

Our method integrated a plug-and-play (PnP) diffusion prior with an ℓ_q -norm fidelity term to address general noise. As a generalizable PnP framework, its core innovation lies in

how the denoiser (induced by a generative model) is incorporated into the iterative solver. Most PnP methods differ primarily in their choice of denoiser. Conventional denoisers are typically trained for Gaussian noise at a fixed noise level, requiring paired datasets. In contrast, diffusion models naturally yield a continuum of denoisers across a range of noise levels. To exploit this, we introduce a renoising step that ensures the input to the denoiser matches the Gaussian noise distribution expected by the diffusion model.

Due to the modularity of PnP models, all available generative models can be integrated in our proposed framework. Hence, given any generative model, including the flow matching model (Lipman et al. 2023), diffusion model (Ho, Jain, and Abbeel 2020), and rectified flow model (Liu, Gong et al. 2023), we can follow the algorithmic flowchart to restore images from its degradations. The minor modification is the replacement of the denoiser induced by individual models. To show the flexibility of our algorithm, we consider image restoration with general noise removal with integration of diffusion model and flow matching model.

Numerical Experiments

To evaluate the performance of our algorithm for image restoration with general noise via gGSM approximation, we conduct experiments on various image degradations with simulated impulse salt-and-pepper noise. We consider three benchmark datasets with image size 256×256 , including Flickr Faces High Quality (FFHQ) (Karras, Laine, and Aila 2019), ImageNet (Deng et al. 2009), and Animal Faces HQ (AFHQ) with cat category. The publicly available generative models for the first two datasets are trained using diffusion model, and both of them are built on VP SDE. For FFHQ, we use the pretrained DDPM model from (Choi et al. 2021), and for ImageNet, we use the pretrained DDPM model from (Dhariwal and Nichol 2021). For the AFHQ-Cat dataset, we access the pretrained flow matching model (Martin et al. 2024) with a time point definition conversion to follow the convention of flow model. All models are used as-is without task-specific fine-tuning. For FFHQ and ImageNet, we randomly select 100 testing images from 1K tested images to evaluate the algorithm, while for AFHQ, we report the result on 32 testing images. The denoiser of diffusion and flow model can be obtained via Tweedie’s formula.

Evaluated imaging problems The considered four image restoration problems are as follows. (a) Denoising; (b) deblurring using a 61×61 Gaussian kernel; (c) random pixel inpainting with 70% missing pixels. (d) super-resolution with $4 \times$ downsampling via average pooling.

Simulated noisy measurement To evaluate the advantages of ℓ_q -norm loss for simulated impulse noise. We consider corrupting the measurement with salt-and-pepper noise. We consider two noise levels $r_{sp} = 0.5, 0.7$, which means that the random 50%, 70% percentage of pixels are corrupted by setting the values to the minimum and maximum values with equality probability.

Quantitative results Within the PnP solving framework, fixing the data fidelity term, we evaluate the effect of different regularizations on the image restorations. The compared regularizations include total variation (TV), DRUnet (Zhang




Figure 3: Visualization of different methods on four tasks with impulse noise for FFHQ and ImageNet.

et al. 2021) pretrained for Gaussian noise removal, and our method with denoiser induced by diffusion and flow model. We also include the original DPS method (Chung et al. 2022) with $q = 2$ as a baseline. All methods share a consistent IRLS-based solver and optimization setup. The stepsize is set as $s_i = 1/\|\mathbf{A}^T(\mathbf{W}_i(\mathbf{A}\mathbf{x}_i - \mathbf{y}))\|_2$ and $T_{inter} = 1$. For TV regularization, we run 60-step proximal minimiza-

tion and the trade-off parameter is set to $s_i\eta_t$. For DRUnet, the noise level is set to $s_i\eta_t$ as well. For all experiments, we set $q = 0.5$. We evaluate the results using PSNR \uparrow and SSIM \uparrow to quantify fidelity to the original images.

Table 1 presents quantitative results under $r_{sp} = 0.5$ salt-and-pepper noise. Across all datasets and tasks, our method consistently outperforms all competing approaches. Visual

comparisons in Figures 3 and 4 highlight the superior reconstruction quality, especially for challenging tasks like inpainting and super-resolution. DPS achieves visually detailed outputs but suffers from measurement inconsistency, often producing a gray layer over the image. These results underscore the effectiveness of the ℓ_q -norm data fidelity in handling impulse noise and the superiority of generative priors (diffusion and flow) over hand-crafted or conventional denoisers. Additional qualitative examples are provided in the Appendix of the extended arXiv version.

Method	Denoise		Deblur		Inpaint		SR		
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	
FFHQ	PnP-TV	25.22	0.763	20.39	0.633	22.74	0.686	16.82	0.552
	DRUnet	26.42	0.765	25.10	0.719	16.33	0.572	12.90	0.464
	DPS	16.71	0.565	16.36	0.549	11.25	0.416	10.83	0.065
	Ours	28.36	0.828	25.45	0.729	23.59	0.727	20.24	0.621
ImageNet	PnP-TV	22.84	0.683	18.20	0.491	21.70	0.591	14.13	0.436
	DRUnet	24.63	0.675	22.68	0.589	17.08	0.506	10.45	0.237
	DPS	16.36	0.412	12.32	0.215	7.96	0.248	8.97	0.037
	Ours	25.81	0.679	23.33	0.613	21.71	0.593	16.36	0.431
AFHQ	PnP-TV	24.59	0.729	20.16	0.560	23.21	0.643	16.98	0.475
	DRUnet	27.43	0.770	24.74	0.654	16.12	0.505	14.00	0.434
	Ours	29.73	0.879	25.26	0.677	26.16	0.761	20.94	0.581

Table 1: Quantitative results of different methods for different tasks on three datasets with 0.5 salt-and-pepper noise.

Ablation study on the choice of loss We conduct an ablation study using AFHQ dataset to investigate the impact of the choice of q in the ℓ_q -norm data fidelity term on restoration performance. Our iteratively reweighted least-squares solver is applicable for any $q \in (0, 2]$. As shown in Introduction, for impulse noise removal, setting q in interval $(0, 1]$ is better. See Table 2 for the final PSNR/SSIM results with different q values. Smaller q values (e.g., $q \in (0, 0.5]$) consistently yield better performance, and the performance degrades as q increases beyond 0.5. See Figure 5 for the illustration of the effect of q values.

q	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
PSNR	29.62	29.69	29.72	29.73	29.66	29.49	29.16	28.59	27.76
SSIM	0.871	0.876	0.879	0.879	0.863	0.852	0.838	0.82	0.799

Table 2: Performance comparison of image denoising on AFHQ-Cat dataset using different q values.

Conclusion

Within the general play-and-play (PnP) framework for image restoration, we propose using ℓ_q norm fidelity term to address general noise. Leveraging the strong diffusion prior as a regularization, we utilize the iteratively reweighted least-squares based method to efficiently solve the diffusion prior regularized ℓ_q norm minimization problem. The

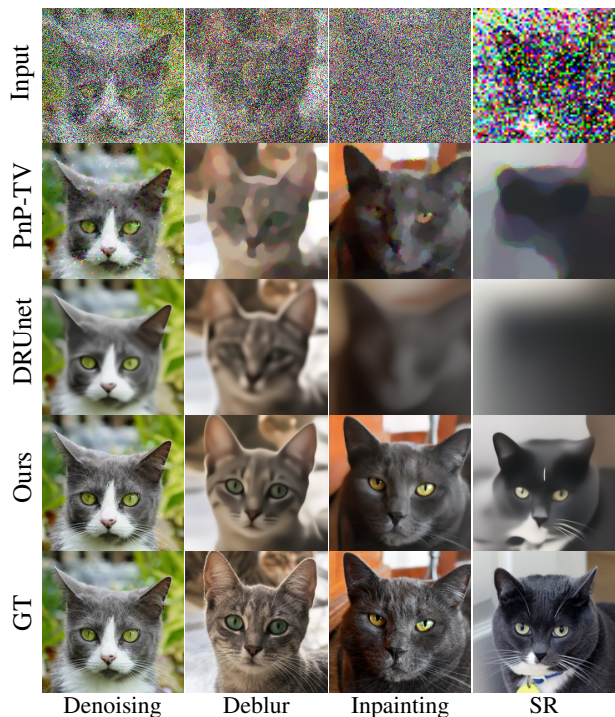


Figure 4: Visualization of different methods on four tasks with impulse noise for AFHQ-Cat.

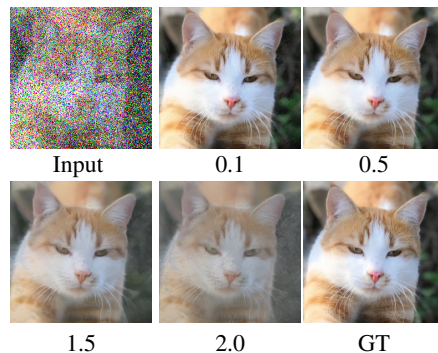


Figure 5: Visualization of restorations from different q values for image denoising on AFHQ-Cat.

core subproblem is tackled via a forward-backward splitting method, where the proximal step involving the diffusion prior is approximated through a denoising operation. This denoising is efficiently performed using the induced denoiser of the pretrained diffusion model. Experiments are conducted to evaluate our proposed method on three datasets with four imaging tasks, utilizing both diffusion and flow-based generative models as priors. Quantitative results and visualizations showcase the superior advantages of our method with the integration of plug-and-play diffusion prior and ℓ_q norm loss over other TV regularizer and deep denoiser. Our method also outperforms the original DPS with $q = 2$, as it is tailored for Gaussian noise.

Acknowledgments

Ji Li was supported by the National Natural Science Foundation of China (Grant No. 12571472) and supported by the Open Project of Key Laboratory of Mathematics and Information Networks (Beijing University of Posts and Telecommunications), Ministry of Education, China, under Grant No. KF202401). Chao Wang was supported by the National Natural Science Foundation of China (Grant No. 12526572).

References

- Adler, J.; and Öktem, O. 2018. Learned primal-dual reconstruction. *IEEE transactions on medical imaging*, 37(6): 1322–1332.
- Aggarwal, H. K.; Mani, M. P.; and Jacob, M. 2018. MoDL: Model-based deep learning architecture for inverse problems. *IEEE transactions on medical imaging*, 38(2): 394–405.
- Beck, A.; and Teboulle, M. 2009. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE transactions on image processing*, 18(11): 2419–2434.
- Blomgren, P.; Chan, T. F.; Mulet, P.; and Wong, C.-K. 1997. Total variation image restoration: numerical methods and extensions. In *Proceedings of international conference on image processing*, volume 3, 384–387. IEEE.
- Chambolle, A.; Caselles, V.; Cremers, D.; Novaga, M.; and Pock, T. 2010. An introduction to total variation for image analysis. *Theoretical foundations and numerical methods for sparse recovery*, 9(263-340): 227.
- Chan, T.; Esedoglu, S.; Park, F.; and Yip, A. 2006. Total variation image restoration: Overview and recent developments. *Handbook of mathematical models in computer vision*, 17–31.
- Chartrand, R.; and Yin, W. 2008. Iteratively reweighted algorithms for compressive sensing. In *2008 IEEE international conference on acoustics, speech and signal processing*, 3869–3872. IEEE.
- Choi, J.; Kim, S.; Jeong, Y.; Gwon, Y.; and Yoon, S. 2021. ILVR: Conditioning Method for Denoising Diffusion Probabilistic Models. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 14347–14356. IEEE.
- Chung, H.; Kim, J.; Mccann, M. T.; Klasky, M. L.; and Ye, J. C. 2022. Diffusion Posterior Sampling for General Noisy Inverse Problems. In *The Eleventh International Conference on Learning Representations*.
- Chung, H.; Lee, S.; and Ye, J. C. 2023. Fast Diffusion Sampler for Inverse Problems by Geometric Decomposition. *arXiv preprint arXiv:2303.05754*.
- Dabov, K.; Foi, A.; Katkovnik, V.; and Egiazarian, K. 2007. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8): 2080–2095.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. volume 34, 8780–8794.
- Eksioglu, E. 2016. Decoupled algorithm for MRI reconstruction using nonlocal block matching model: BM3D-MRI. *Journal of Mathematical Imaging and Vision*, 56(3): 430–440.
- Figueiredo, M. A.; and Nowak, R. D. 2001. Wavelet-based image estimation: An empirical Bayes approach using Jeffrey’s noninformative prior. *IEEE Transactions on Image Processing*, 10(9): 1322–1331.
- Gilton, D.; Ongie, G.; and Willett, R. 2019. Neumann networks for inverse problems in imaging. *arXiv preprint arXiv:1901.03707*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. volume 33, 6840–6851.
- Jin, K. H.; McCann, M. T.; Froustey, E.; and Unser, M. 2017. Deep convolutional neural network for inverse problems in imaging. *IEEE transactions on image processing*, 26(9): 4509–4522.
- Kamilov, U. S.; Mansour, H.; and Wohlberg, B. 2017. A plug-and-play priors approach for solving nonlinear imaging inverse problems. *IEEE Signal Processing Letters*, 24(12): 1872–1876.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Koshelev, I.; and Lefkimmatis, S. 2023. Iterative Reweighted Least Squares Networks With Convergence Guarantees for Solving Inverse Imaging Problems. *arXiv preprint arXiv:2308.05745*.
- Li, J.; and Wang, C. 2024. DiffFPR: Diffusion prior for over-sampled Fourier phase retrieval. In *Forty-first International Conference on Machine Learning*.
- Lipman, Y.; Chen, R. T.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. In *The Eleventh International Conference on Learning Representations*.
- Liu, X.; Gong, C.; et al. 2023. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. In *The Eleventh International Conference on Learning Representations*.
- Mardani, M.; Song, J.; Kautz, J.; and Vahdat, A. 2024. A Variational Perspective on Solving Inverse Problems with Diffusion Models. In *The Twelfth International Conference on Learning Representations*.
- Martin, S.; Gagneux, A.; Hagemann, P.; and Steidl, G. 2024. PnP-Flow: Plug-and-Play Image Restoration with Flow Matching. *arXiv preprint arXiv:2410.02423*.
- Meinhardt, T.; Moller, M.; Hazirbas, C.; and Cremers, D. 2017. Learning proximal operators: Using denoising networks for regularizing inverse imaging problems. In *ICCV*, 1781–1790.
- Metzler, C. A.; Maleki, A.; and Baraniuk, R. 2015. BM3D-AMP: A new image recovery algorithm based on BM3D denoising. In *ICIP*, 3116–3120. IEEE.

- Metzler, C. A.; Maleki, A.; and Baraniuk, R. G. 2016. BM3D-PRGAMP: Compressive phase retrieval based on BM3D denoising. In *2016 IEEE International Conference on Image Processing (ICIP)*, 2504–2508. IEEE.
- Ochs, P.; Dosovitskiy, A.; Brox, T.; and Pock, T. 2015. On iteratively reweighted algorithms for nonsmooth nonconvex optimization in computer vision. *SIAM Journal on Imaging Sciences*, 8(1): 331–372.
- Parikh, N.; Boyd, S.; et al. 2014. Proximal algorithms. *Foundations and trends® in Optimization*, 1(3): 127–239.
- Reehorst, E. T.; and Schniter, P. 2018. Regularization by denoising: Clarifications and new interpretations. *IEEE transactions on computational imaging*, 5(1): 52–67.
- Romano, Y.; Elad, M.; and Milanfar, P. 2017. The little engine that could: Regularization by denoising (RED). *SIAM Journal on Imaging Sciences*, 10(4): 1804–1844.
- Rudin, L. I.; Osher, S.; and Fatemi, E. 1992. Nonlinear total variation based noise removal algorithms. *Physica D: non-linear phenomena*, 60(1-4): 259–268.
- Ryu, E.; Liu, J.; Wang, S.; Chen, X.; Wang, Z.; and Yin, W. 2019. Plug-and-play methods provably converge with properly trained denoisers. In *International Conference on Machine Learning*, 5546–5557. PMLR.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *ICLR*.
- Song, J.; Vahdat, A.; Mardani, M.; and Kautz, J. 2022. Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations*.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. volume 32.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Sun, Y.; Babu, P.; and Palomar, D. P. 2016. Majorization-minimization algorithms in signal processing, communications, and machine learning. *IEEE Transactions on Signal Processing*, 65(3): 794–816.
- Venkatakrishnan, S. V.; Bouman, C. A.; and Wohlberg, B. 2013. Plug-and-play priors for model based reconstruction. In *2013 IEEE Global Conference on Signal and Information Processing*, 945–948. IEEE.
- Vogel, C. R.; and Oman, M. E. 1996. Iterative methods for total variation denoising. *SIAM Journal on Scientific Computing*, 17(1): 227–238.
- Wang, Y.; Yu, J.; and Zhang, J. 2022. Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model. In *The Eleventh International Conference on Learning Representations*.
- Windygga, P. S. 2001. Fast impulsive noise removal. *IEEE transactions on image processing*, 10(1): 173–179.
- Zhang, K.; Li, Y.; Zuo, W.; Zhang, L.; Van Gool, L.; and Timofte, R. 2020. Plug-and-play image restoration with deep denoiser prior. *arXiv preprint arXiv:2008.13751*.
- Zhang, K.; Li, Y.; Zuo, W.; Zhang, L.; Van Gool, L.; and Timofte, R. 2021. Plug-and-Play Image Restoration with Deep Denoiser Prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10): 6360–6376.
- Zhang, K.; Zuo, W.; Chen, Y.; Meng, D.; and Zhang, L. 2017. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7): 3142–3155.
- Zhang, L.; Wei, W.; Zhang, Y.; Tian, C.; and Li, F. 2015. Reweighted laplace prior based hyperspectral compressive sensing for unknown sparsity. In *CVPR*, 2274–2281.
- Zhu, Y.; Zhang, K.; Liang, J.; Cao, J.; Wen, B.; Timofte, R.; and Van Gool, L. 2023. Denoising Diffusion Models for Plug-and-Play Image Restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 1219–1229.