

Zero-Shot Adaptation for Approximate Posterior Sampling of Diffusion Models in Inverse Problems

Yaşar Utku Alçalar[✉] and Mehmet Akçakaya[✉]

University of Minnesota, Minneapolis
{alcal029, akcakaya}@umn.edu

Abstract. Diffusion models have emerged as powerful generative techniques for solving inverse problems. Despite their success in a variety of inverse problems in imaging, these models require many steps to converge, leading to slow inference time. Recently, there has been a trend in diffusion models for employing sophisticated noise schedules that involve more frequent iterations of timesteps at lower noise levels, thereby improving image generation and convergence speed. However, application of these ideas for solving inverse problems with diffusion models remain challenging, as these noise schedules do not perform well when using empirical tuning for the forward model log-likelihood term weights. To tackle these challenges, we propose zero-shot approximate posterior sampling (ZAPS) that leverages connections to zero-shot physics-driven deep learning. ZAPS fixes the number of sampling steps, and uses zero-shot training with a physics-guided loss function to learn log-likelihood weights at each irregular timestep. We apply ZAPS to the recently proposed diffusion posterior sampling method as baseline, though ZAPS can also be used with other posterior sampling diffusion models. We further approximate the Hessian of the logarithm of the prior using a diagonalization approach with learnable diagonal entries for computational efficiency. These parameters are optimized over a fixed number of epochs with a given computational budget. Our results for various noisy inverse problems, including Gaussian and motion deblurring, inpainting, and super-resolution show that ZAPS reduces inference time, provides robustness to irregular noise schedules and improves reconstruction quality. Code is available at <https://github.com/ualcalar17/ZAPS>.

Keywords: Diffusion Models · Zero-Shot Learning · Inverse Problems · Plug-and-Play (PnP) Methods · Unrolled Networks · Bayesian Methods

1 Introduction

The forefront of deep generative models is now dominated by diffusion models [16, 28, 30, 32, 34] in the intricate task of image generation [11]. Their capabilities extend across various domains, including computer vision [2], natural language processing [17] and temporal data modeling [1]. Recently, diffusion models also showed great success in solving noiseless [5, 7, 33, 34] and noisy inverse problems [6, 21, 29, 31], owing to their capability to model complicated



Fig. 1: Representative results of our algorithm for four distinct noisy inverse problems ($\sigma = 0.05$), showing the ground truth (GT), measurement and reconstruction.

high-dimensional distributions. Linear inverse problems utilize a known forward model given by

$$\mathbf{y} = \mathbf{A}\mathbf{x}_0 + \mathbf{n},$$

and aim to deduce the underlying signal/image $\mathbf{x}_0 \in \mathbb{R}^n$ from measurements $\mathbf{y} \in \mathbb{R}^m$, where $\mathbf{n} \in \mathbb{R}^m$ is measurement noise. In practical situations, the forward operator $\mathbf{A} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is either incomplete or ill-conditioned, necessitating the use of prior information about the signal. Posterior sampling approaches use diffusion models as generative priors and incorporates information from both the data distribution and the forward physics model, allowing for sampling from the posterior distribution $p(\mathbf{x}|\mathbf{y})$ using the given measurement \mathbf{y} [21]. In this context, using Bayes' rule, $p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{x})p(\mathbf{y}|\mathbf{x})}{p(\mathbf{y})}$, the problem-specific score is

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}|\mathbf{y}) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}), \quad (1)$$

where $\nabla_{\mathbf{x}_t} \log p(\mathbf{x})$ is approximated via the learned score model $s_\theta(\mathbf{x}_t, t)$. Many of these strategies utilize a plug-and-play (PnP) approach, using a pre-trained unconditional diffusion model as a prior [4, 9, 13, 18, 24, 37], and integrate the forward model during inference to address various inverse problem tasks.

The complexity for these approaches arises in obtaining the latter forward model log-likelihood term in Eq. (1), which guides the diffusion to a target

class [11, 28]. While exact calculation is intractable, several approaches have been proposed to approximate this term. Among these, RED-diff [25] employs a variational sampler that uses a combination of measurement consistency loss and score matching regularization. Another technique, DSG [43], uses a spherical Gaussian constraint for denoising steps, allowing for larger step sizes. A class of methods utilize projections onto the convex measurement subspace after the unconditional update through score model [5, 8, 34]. Although these projections improve consistency between measurements and the sample, they are noted to lead to artifacts, such as boundary effects [7]. Thus, more recent approaches aimed to approximate the log-likelihood term in Eq. (1) different ways. Noting

$$p_t(\mathbf{y}|\mathbf{x}_t) = \int_{\mathbf{x}_0} p(\mathbf{x}_0|\mathbf{x}_t)p(\mathbf{y}|\mathbf{x}_0)d\mathbf{x}_0, \quad (2)$$

DPS [6] uses the posterior mean $\hat{\mathbf{x}}_0 = \hat{\mathbf{x}}_0(\mathbf{x}_t) \triangleq \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t] = \mathbb{E}_{\mathbf{x}_0 \sim p(\mathbf{x}_0|\mathbf{x}_t)}[\mathbf{x}_0]$, to approximate $p(\mathbf{y}|\mathbf{x}_t) = \mathbb{E}_{\mathbf{x}_0 \sim p(\mathbf{x}_0|\mathbf{x}_t)}[p(\mathbf{y}|\mathbf{x}_0)]$ as

$$p(\mathbf{y}|\mathbf{x}_t) = \mathbb{E}_{\mathbf{x}_0 \sim p(\mathbf{x}_0|\mathbf{x}_t)}[p(\mathbf{y}|\mathbf{x}_0)] \simeq p\left(\mathbf{y}|\mathbb{E}_{\mathbf{x}_0 \sim p(\mathbf{x}_0|\mathbf{x}_t)}[\mathbf{x}_0]\right) = p(\mathbf{y}|\hat{\mathbf{x}}_0).$$

Another technique, IIGDM [31] approximates Eq. (2) as a Gaussian centered around $\mathbf{A}\hat{\mathbf{x}}_0$

$$\int_{\mathbf{x}_0} p(\mathbf{x}_0|\mathbf{x}_t)p(\mathbf{y}|\mathbf{x}_0)d\mathbf{x}_0 \simeq \mathcal{N}(\mathbf{A}\hat{\mathbf{x}}_0, r_t^2\mathbf{A}\mathbf{A}^\top + \sigma_y^2\mathbf{I}), \quad (3)$$

and uses it for guidance. In these works, log-likelihood weights (or gradient step sizes), $\{\zeta_t\}$ are introduced to further control the reconstruction as

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}|\mathbf{y}) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}) + \zeta_t \nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}). \quad (4)$$

While DPS demonstrates high performance in various inverse problem tasks, it suffers from the drawback of requiring a large number of sampling steps, resulting in prolonged reconstruction time. IIGDM accelerates this process by adopting regular (linear) jumps approach across the schedule. However, utilizing more complicated schedules, where the jumps are irregular introduces a challenge, as it requires distinct log-likelihood weights, ζ_t , for each timestep. Heuristic adjustment of these weights is difficult and frequently leads to undesirable outcomes. In this work, by taking an inspiration from zero-shot/test-time self-supervised models [35, 42] we propose to learn the log-likelihood weights for a fixed number of sampling steps and fine-tune them over a few epochs. It is crucial to note that fine-tuning DPS (or IIGDM) entails saving computational graphs for each unroll, leading to memory issues and slow backpropagation. Thus, we also propose to approximate the Hessian of the data probability using a wavelet-based diagonalization strategy [12], and learn these diagonal values for each timestep as well. Fig. 1 shows representative results for our method. Our key contributions include:

- We introduce zero-shot approximate posterior sampling (ZAPS), leveraging zero-shot learning for dynamic automated hyperparameter tuning in the inference phase to improve solution of noisy inverse problems via diffusion

models. This method fortifies the robustness of the sampling process, attaining a state-of-the-art performance [6, 21, 31] in sampling outcomes. To the best of our knowledge, our method is the first attempt to learn the log-likelihood weights for solving inverse problems via diffusion models by using a measurement-consistent loss when the sampling noise schedule consists of irregular jumps across timesteps.

- We provide a well-designed approximation for the Hessian of the logarithm of the prior, enabling a computationally efficient and trainable posterior computation.
- We showcase the efficacy of incorporating a learnable log-likelihood weights for each diffusion step during the reverse diffusion process through both quantitative and qualitative assessments on FFHQ and ImageNet datasets. Our approach not only outperforms state-of-the-art, but it also substantially reduces the required number of sampling steps from 1000 to ~ 20 -to-30, facilitating convergence with fewer total neural function evaluations (NFEs).

2 Related Works

Diffusion Models. During training, diffusion models [16, 34] add Gaussian noise to an image with a fixed increasing variance schedule, e.g. linear or exponential, $\beta_1, \beta_2, \dots, \beta_T$ until pure noise is obtained, and learns a reverse diffusion process, where a neural network is trained to gradually remove noise and reconstruct the original image. Let $\mathbf{x}_0 \sim p_{\text{data}}(x)$ be samples from the data distribution, and $\mathbf{x}_{\{1:T\}} \in \mathbb{R}^d$ be noisy latent variables. By taking $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, the Markovian forward process can be written as

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t|\sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I}). \quad (5)$$

By using the reparameterization trick and Eq. (5), \mathbf{x}_t can be sampled as

$$\mathbf{x}_t(\mathbf{x}_0, \boldsymbol{\epsilon}) = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon} \quad \text{where } \boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\epsilon}; 0, \mathbf{I}). \quad (6)$$

Consequently, denoising diffusion probabilistic models (DDPMs) [16] learns the reverse process by minimizing a lower bound on the log prior via:

$$L_t(\theta) = \mathbb{E}_{t, \mathbf{x}_0, \boldsymbol{\epsilon}} \|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\mathbf{x}_t(\mathbf{x}_0, \boldsymbol{\epsilon}), t)\|_2^2. \quad (7)$$

Furhtermore, it can be shown that epsilon matching in Eq. (7) is analogous to the denoising score matching (DSM) [32, 39] objective up to a constant:

$$\min_{\theta} \mathbb{E}_{\mathbf{x}_t, \mathbf{x}_0, \boldsymbol{\epsilon}} \|\mathbf{s}_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t|\mathbf{x}_0)\|_2^2, \quad (8)$$

in which $\mathbf{s}_\theta(\mathbf{x}_t, t) = -\frac{\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)}{\sqrt{1 - \bar{\alpha}_t}}$. Using Tweedie’s formula and Eq. (6), posterior mean for $p(\mathbf{x}_0|\mathbf{x}_t)$ can be found as:

$$\hat{\mathbf{x}}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{x}_t + (1 - \bar{\alpha}_t)\mathbf{s}_\theta(\mathbf{x}_t, t) \right). \quad (9)$$

Sampling \mathbf{x}_{t+1} from $p(\mathbf{x}_{t+1}|\mathbf{x}_t)$ can be done using ancestral sampling by iteratively computing:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_{t-1} - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z}, \quad (10)$$

where $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$ and $\sigma_t^2 = \tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t$. It is also worth noting that the DDPM is equivalent to the variance preserving stochastic differential equations (VP-SDEs) [34].

Solving Inverse Problems via Diffusion Models. When solving inverse problems via diffusion models, the main challenge is to find an approximation to the log-likelihood term, $\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x})$, as discussed earlier. One recent method, denoising diffusion restoration models (DDRM) [21], utilizes a spectral domain approach, allowing the incorporation of noise from the measurement domain into the spectral domain through singular value decomposition (SVD). However, the application of SVD is computationally expensive [6]. Manifold Constrained Gradient (MCG) [7] method applies projections after the MCG correction as:

$$\mathbf{x}'_{t-1} = f(\mathbf{x}_t, \mathbf{s}_\theta) - \zeta \nabla_{\mathbf{x}_t} \|\mathbf{K}(\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0)\|_2^2 + g(\mathbf{x}_t)\mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(0, \mathbf{I}), \quad (11)$$

$$\mathbf{x}_{t-1} = \mathbf{H}\mathbf{x}'_{t-1} + \mathbf{b}, \quad (12)$$

where ζ and \mathbf{H} are dependent on noise covariance. MCG update of Eq. (11) projects estimates onto the measurement subspace, thus they may fall off from the data manifold [6]. Hence, DPS proposes to update without projections as:

$$\mathbf{x}_{t-1} = \mathbf{x}'_{t-1} - \zeta_t \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0\|_2^2, \quad (13)$$

Note Eq. (13) is equivalent to Eq. (11) when $\mathbf{K} = \mathbf{I}$, and it reduces to the following when the forward operator is linear:

$$\mathbf{x}_{t-1} = \mathbf{x}'_{t-1} + \zeta_t \frac{\partial \hat{\mathbf{x}}_0}{\partial \mathbf{x}_t} \mathbf{A}^\top (\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0) \quad (14)$$

PIGDM [31], on the other hand, utilizes a Gaussian centered around $\hat{\mathbf{x}}_0$ that is defined in Eq. (9) to obtain the following score approximation:

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t) \simeq \frac{\partial \hat{\mathbf{x}}_0}{\partial \mathbf{x}_t} \mathbf{A}^\top (r_t^2 \mathbf{A}\mathbf{A}^\top + \sigma_y^2 \mathbf{I})^{-1} (\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0). \quad (15)$$

In cases where there is no measurement noise ($\sigma_y = 0$), Eq. (15) simplifies to:

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t) \simeq r_t^{-2} \frac{\partial \hat{\mathbf{x}}_0}{\partial \mathbf{x}_t} \mathbf{A}^\dagger (\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0) \quad (16)$$

where \mathbf{A}^\dagger denotes the Moore-Penrose pseudoinverse of \mathbf{A} . We note that using Woodbury matrix identity (derived in SuppMat), one can simplify Eq. (15) to:

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t) \simeq \frac{\partial \hat{\mathbf{x}}_0}{\partial \mathbf{x}_t} (\mathbf{A}^\top \mathbf{A} + \eta \mathbf{I})^{-1} \mathbf{A}^\top (\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0), \quad \text{where } \eta = \frac{\sigma_y^2}{r_t^2}. \quad (17)$$

From Eq. (17), the similarity between DPS and IIGDM updates can be seen, with $(\mathbf{A}^\top \mathbf{A} + \eta \mathbf{I})^{-1}$ term being the difference. Note the DPS update in Eq. (13) works with non-linear operators, while IIGDM’s update does not rely on the differentiability of the forward operator, as long as a pseudo-inverse-like operation can be derived.

Improved Irregular Noise Schedules for Image Generation. Diffusion models typically utilize well-defined fixed noise schedules, with examples including linear or exponential ones. Lately, more sophisticated methods have been developed that sweep across these schedules and take samples in irregular timesteps [11,19] for unconditional image generation. The idea behind this strategy hinges on more frequent sampling for lower noise levels, making it possible to use considerably less number of sampling steps.

Most of the aforementioned studies that solve inverse problems via diffusion models used the same number of steps that the unconditional diffusion model was trained for [6,7,34]. Nonetheless, there has been a notable trend favoring shorter schedules characterized by linear jumps for inverse problems, where the log-likelihood weights were hand-tuned by trial-and-error [25,31] when using reduced number of steps. While these approaches have proven effective, they still require a large number of sampling steps or heuristic tuning of the log-likelihood weights, $\{\zeta_t\}$ in Eq. (4) to achieve good performance. The former issue leads to lengthy and potentially impractical computational times, while the latter issue results in generalizability difficulties for adoption at different measurement noise levels and variations in the measurement operators. Furthermore, the irregular jump strategy that has been powerful for image generation has not garnered significant attention for inverse problems, mainly due to the impracticality of empirically tuning the log-likelihood weights. Thus, a method that automatically selects and adjusts log-likelihood weights based on the provided measurements for arbitrary noise schedules, instead of requiring manual tuning, holds significant potential for improving robustness and image quality.

3 Methodology

3.1 Zero-shot Fine Tuning of Log-Likelihood Weights

In this work, we propose a robust automated approach for setting the log-likelihood weights at each timestep for arbitrary noise sampling schedules to improve posterior sampling with the given measurements during inference. This allows for a stable reconstruction for different sweeps across noise schedules. Furthermore, the weights themselves are image-specific, which improves the performance compared to the former approaches. For estimating the likelihood in Eq. (1), we use the update in DPS [6]:

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t) \simeq \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0\|_2^2 = -\frac{\partial \hat{\mathbf{x}}_0}{\partial \mathbf{x}_t} \mathbf{A}^\top (\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0), \quad (18)$$

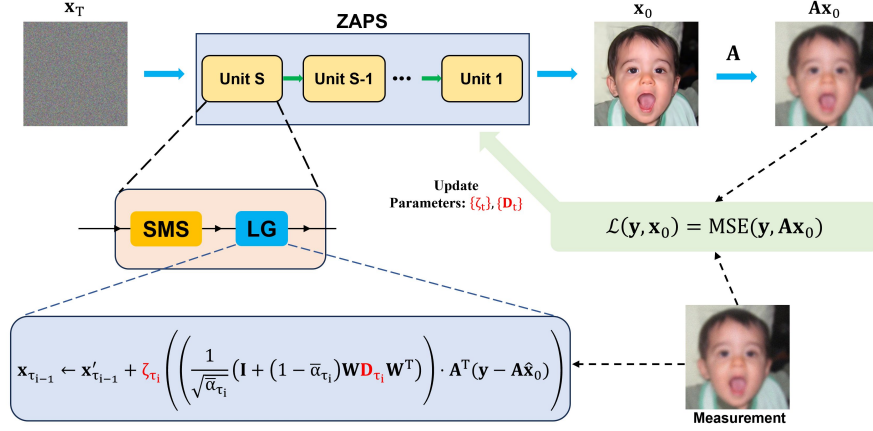


Fig. 2: Our zero-shot approximate posterior sampling (ZAPS) approach unrolls the sampling process for a fixed number of S steps for arbitrary/irregular noise schedules, alternating between score model sampling (SMS) and likelihood guidance (LG). Our zero-shot fine-tuning approach has two key components: 1) The Hessian of the log prior is approximated using a discrete wavelet transform diagonalization technique, 2) Both the diagonal matrices, $\{\mathbf{D}_t\}$ and the log-likelihood weights, $\{\zeta_t\}$ are updated during fine-tuning. The fine-tuning is done for a fixed number of epochs with a given NFE budget, yielding a faster and more robust *adaptive* inverse problem solver.

although as noted before, the Π GDM [31] update in Eq. (17) is also similar. Thus we emphasize that while we chose DPS as baseline for its versatility in inverse problems, our ZAPS strategy is applicable to other diffusion models for inverse problems. Recalling the definition of $\hat{\mathbf{x}}_0$ in Eq. (9), we note

$$\frac{\partial \hat{\mathbf{x}}_0}{\partial \mathbf{x}_t} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{I} + (1 - \bar{\alpha}_t) \frac{\partial \mathbf{s}_\theta(\mathbf{x}_t, t)}{\partial \mathbf{x}_t} \right). \quad (19)$$

Thus, ignoring the calculation and storage of the matrix $\frac{\partial \mathbf{s}_\theta(\mathbf{x}_t, t)}{\partial \mathbf{x}_t}$ for now, one needs to fine tune the log-likelihood weights $\{\zeta_t\}$ in

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}) + \zeta_t \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{I} + (1 - \bar{\alpha}_t) \frac{\partial \mathbf{s}_\theta(\mathbf{x}_t, t)}{\partial \mathbf{x}_t} \right) \mathbf{A}^\top (\mathbf{y} - \mathbf{A} \hat{\mathbf{x}}_0). \quad (20)$$

This is done based on the concept of algorithm unrolling [14, 15, 22] in physics-driven deep learning by fixing the number of sampling steps T . Then the whole posterior sampling process is described as alternating between DDPM sampling using the pre-trained unconditional score model, followed by the log-likelihood term guidance in Eq. (20) for T steps. This “unrolled” network is fine-tuned end-to-end, where the only updates are made to $\{\zeta_t\}$ and no fine-tuning is performed on the unconditional score function, $\mathbf{s}_\theta(\mathbf{x}_t, t)$. This also alleviates the need for backpropagation across the score function network, leading to further savings in computational time. The fine-tuning is performed using a physics-inspired loss

Algorithm 1 ZAPS: Zero-Shot Approximate Posterior Sampling

Require: T , \mathbf{y} , $\{\tilde{\sigma}_i\}_{i=1}^T$, orthogonal DWT (\mathbf{W})

- 1: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 2: $\tau \subset [1, \dots, T]$ extending over a length of $S < T$
- 3: **for** epoch **in** range(epochs) **do**
- 4: **for** $i = S, \dots, 1$ **do**
- 5: $\hat{\mathbf{s}} \leftarrow \mathbf{s}_\theta(\mathbf{x}_{\tau_i}, \tau_i)$ \triangleright *Score computation*
- 6: $\hat{\mathbf{x}}_0 \leftarrow \frac{1}{\sqrt{\bar{\alpha}_{\tau_i}}}(\mathbf{x}_{\tau_i} + (1 - \bar{\alpha}_{\tau_i})\hat{\mathbf{s}})$ \triangleright *Tweedie denoising*
- 7: $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $\tau_i > 1$, else $\mathbf{z} = \mathbf{0}$
- 8: $\mathbf{x}'_{\tau_{i-1}} \leftarrow \frac{\sqrt{\bar{\alpha}_{\tau_i}(1-\bar{\alpha}_{\tau_{i-1}})}}{1-\bar{\alpha}_{\tau_i}}\mathbf{x}_{\tau_i} + \frac{\sqrt{\bar{\alpha}_{\tau_{i-1}}\beta_{\tau_i}}}{1-\bar{\alpha}_{\tau_i}}\hat{\mathbf{x}}_0 + \tilde{\sigma}_{\tau_i}\mathbf{z}$
- 9: $\mathbf{x}_{\tau_{i-1}} \leftarrow \mathbf{x}'_{\tau_{i-1}} + \zeta_{\tau_i} \left(\left(\frac{1}{\sqrt{\bar{\alpha}_{\tau_i}}} \left(\mathbf{I} + (1 - \bar{\alpha}_{\tau_i}) \mathbf{W} \mathbf{D}_{\tau_i} \mathbf{W}^\top \right) \right) \cdot \mathbf{A}^\top (\mathbf{y} - \mathbf{A} \hat{\mathbf{x}}_0) \right)$
- 10: **end for**
- 11: Update network parameters $\{\zeta_t\}$ and $\{\mathbf{D}_t\}$
- 12: **end for**
- 13: **return** \mathbf{x}_0

function that evaluates the consistency of the final estimate and the measurements: $\mathcal{L}(\mathbf{y}, \mathbf{x}_0) = \|\mathbf{y} - \mathbf{A}\mathbf{x}_0\|_2^2$. High-level explanation for our algorithm is given in Fig. 2.

3.2 Approximation for the Hessian of the Log Prior

Implementing the zero-shot update for Eq. (20) poses various challenges, since backpropagation through the unrolled network to update all $\{\zeta_t\}$ requires another backpropagation through the Jacobian of the score function at each time step. This can only be done by retaining the computational graphs that are created when calculating the Jacobian term in Eq. (20), which quickly explodes memory requirements, especially when the number of sampling steps increases. Also, backpropagating through multiple graphs at the end to only update the log-likelihood weights is time-inefficient and causes prolonged sampling times. Hence, we propose to approximate the Jacobian using inspirations from wavelet-based signal processing techniques and propose to learn this approximation to improve the overall outcome. Noting that $\mathbf{s}_\theta(\mathbf{x}_t, t)$ in Eq. (19) is an approximation of the log-gradient of the true prior $p(\mathbf{x})$, we have

$$\frac{\partial \hat{\mathbf{x}}_0}{\partial \mathbf{x}_t} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{I} + (1 - \bar{\alpha}_t) \frac{\partial^2 \log p_t(\mathbf{x}_t)}{\partial \mathbf{x}_t^2} \right). \quad (21)$$

In order to make a backpropagation to update these weights, one needs to calculate the Hessian matrix, $\frac{\partial^2 \log p_t(\mathbf{x}_t)}{\partial \mathbf{x}_t^2}$ given in Eq. (21). This matrix is the negative of the observed Fisher information matrix, whose expected value is the Fisher information matrix. It is also known that in the limit, it approximates

the inverse covariance matrix of the maximum likelihood estimator. Furthermore, under mild assumptions about continuity of the prior, the observed Fisher information matrix is symmetric. Thus, an appropriate decorrelating unitary matrix can be used to diagonalize it. While finding the desired unitary matrix is equally time-consuming as calculating this Hessian, several pre-determined unitary transforms have been proposed for decorrelation in the signal processing community for different applications [12, 27, 36]. Of particular note is the use of unitary wavelet transforms for Wiener filtering [12], where these transforms were utilized for their tendency to decorrelate data, i.e. approximate the Karhunen-Loeve transform [27]. In this work, we also use these decorrelating properties to approximately diagonalize the Hessian of the log prior, $\frac{\partial^2 \log p_t(\mathbf{x}_t)}{\partial \mathbf{x}_t^2}$ using fixed orthogonal discrete wavelet transforms (DWT):

$$\frac{\partial^2 \log p_t(\mathbf{x}_t)}{\partial \mathbf{x}_t^2} \simeq \mathbf{W} \mathbf{D}_t \mathbf{W}^\top, \quad (22)$$

where \mathbf{W} is an orthogonal DWT. By making this approximation, backpropagation through the score model can also be avoided, and only the diagonal values in distinct $\{\mathbf{D}_t\}$ matrices needs to be learned. Our final algorithm to sample from pure noise with fine-tuning is given in Algorithm 1.

4 Evaluation

4.1 Experimental Setup and Implementation Details

We comprehensively evaluated our method, examining its performance through both qualitative and quantitative analyses using FFHQ [20] and ImageNet [10] datasets with size $256 \times 256 \times 3$. Pre-trained unconditional diffusion models trained on FFHQ and ImageNet were taken from [5] and [11] respectively, and used without retraining. For our experiments, we sampled 1000 images from FFHQ and ImageNet validation sets. All images underwent pre-processing to be normalized in the range $[0, 1]$. During all the evaluations, a Gaussian measurement noise with $\sigma = 0.05$ was used. For the orthogonal DWT, Daubechies 4 wavelet was utilized. For our quantitative evaluations, we employed 30 sampling steps with a schedule of "15,10,5", and 10 epochs for fine-tuning, resulting in a total of 300 NFEs. As noted in [11], superior schedules may exist but it requires substantial computational time to try out all possible schedules. Thus, we opted a schedule that is simple, and samples more frequently at the lower noise levels [11]. More details about the network architectures and hyperparameter choices are given in SuppMat.

4.2 Experiments on Linear Inverse Problems

Problem Setup. We focused on the following linear inverse problems: (1) Gaussian deblurring, (2) inpainting, (3) motion deblurring, (4) super-resolution. For



Fig. 3: Representative images using various methods for solving Gaussian deblurring, motion deblurring and super-resolution ($\times 4$) tasks. Proposed method qualitatively improves upon each method, including the baseline state-of-the-art DPS.

Gaussian deblurring, we considered a kernel of size 61×61 with a standard deviation $\sigma = 3.0$. For inpainting, we considered two different scenarios wherein we randomly masked out 70% and a 128×128 box region of the image, applied uniformly across all three channels. For motion blur, we generated the blur kernel via the code¹, with 61×61 kernel size and 0.5 intensity, as in [6]. Finally, for super-resolution, we considered bicubic downsampling. All measurements are obtained through applying the forward model to the ground truth image.

Comparison Methods. We compared our method with score-SDE [5, 8, 34], manifold constrained gradients (MCG) [7], denoising diffusion restoration models (DDRM) [21], diffusion posterior sampling (DPS) [6] and pseudo-inverse guided diffusion models (IIGDM) [31]. We note that our implementation of score-SDE follows the same strategy as presented in [6]. We referred to the methods that iteratively applied projections onto convex sets (POCS) as score-SDE. Additional comparisons to DDNM [40] and DiffPIR [44] are also provided in SuppMat. All methods were implemented using their respective public repositories.

Quantitative and Qualitative Results. We evaluated our method quantitatively using learned perceptual image patch similarity (LPIPS) distance, structural similarity index (SSIM), and peak signal-to-noise-ratio (PSNR). Representative results in Fig. 3 shows that DDRM yields blurry results in Gaussian deblurring task. DPS improves sharpness across these distinct inverse problem tasks, while ZAPS yields comparable sharpness while exhibiting a higher similarity to the ground truth, all within a third of the total NFEs.

Representative inpainting results in Fig. 4 show that ZAPS substantially improves upon DDRM, a method that uses a slightly lower 20 timesteps, and

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




Fig. 4: Illustrative images using state-of-the-art methods for random (70%) and box (128×128) inpainting. Proposed method improves upon DDRM, while achieving similar performance to IIGDM and DPS, with subtle improvements shown in zoomed insets.

achieves better similarity to the ground truth and sharpness compared to DPS, which uses almost $33\times$ more steps. Similarly, when compared with IIGDM, it is evident that our method gives comparable results even though $3 - 4\times$ fewer number of steps are used. The zoomed insets highlight subtle improvements afforded by our method compared to state-of-the-art DPS and IIGDM, as seen around the eyes.

Tab. 1 and Tab. 2 show the three quantitative metrics for all methods, while Tab. 3 illustrates their computational complexity. ZAPS outperforms Score-SDE, MCG, and our baseline state-of-the-art comparison, DPS, in computational complexity and quantitative performance, yielding faster and improved reconstructions. Although DDRM and IIGDM surpass ZAPS in terms of computational complexity, ZAPS outperforms both methods quantitatively in terms of all three metrics. Furthermore, IIGDM could not be implemented reliably for several lin-

Table 1: Quantitative results for Gaussian deblurring and random inpainting (70%) on FFHQ dataset. Best: **bold**, second-best: underlined. Comparison methods are omitted if they could not be implemented reliably for the given inverse problem task.

Method	Gaussian Deblurring			Random Inpainting		
	LPIPS \downarrow	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	PSNR \uparrow
DPS [6]	<u>0.128</u>	<u>0.718</u>	<u>25.20</u>	0.104	0.811	28.03
MCG [7]	0.558	0.509	15.12	0.145	0.754	25.33
IIGDM [31]	-	-	-	<u>0.086</u>	0.842	26.62
DDRM [21]	0.183	0.702	24.42	0.198	0.741	25.17
Score-SDE [5, 8, 34]	0.571	0.496	15.17	0.224	0.718	24.44
ZAPS (Ours)	0.121	0.757	26.06	0.078	<u>0.813</u>	<u>27.79</u>

Table 2: Quantitative results for motion deblurring and super-resolution ($\times 4$) on FFHQ dataset. Best: **bold**, second-best: underlined. Comparison methods are omitted if they could not be implemented reliably for the given inverse problem task.

Method	Motion Deblurring			Super-Resolution ($\times 4$)		
	LPIPS \downarrow	SSIM \uparrow	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	PSNR \uparrow
DPS [6]	<u>0.143</u>	<u>0.704</u>	<u>24.03</u>	0.168	0.719	23.86
MCG [7]	0.565	0.497	15.10	0.229	0.623	20.74
IIGDM [31]	-	-	-	<u>0.131</u>	<u>0.760</u>	24.48
DDRM [21]	-	-	-	0.175	0.711	<u>24.55</u>
Score-SDE [5, 8, 34]	0.546	0.488	15.02	0.257	0.609	19.13
ZAPS (Ours)	0.141	0.709	24.16	0.104	0.768	26.63

ear inverse problems related to deblurring. We also note that the parameters in ZAPS are adaptive, meaning one can reach the same computational complexity by adjusting total epochs or steps, in trade-off for a slight decrease in performance, as studied in Sec. 4.3.

4.3 Ablation Studies

We conducted three distinct ablation studies to investigate critical aspects of our algorithm’s performance. The first ablation study compared combinations of different timesteps and epochs with a fixed NFE budget, providing a nuanced exploration into the influence of specific combinations on the model’s behavior. Specifically, we explored the reconstruction capabilities of the model qualitatively and quantitatively by varying the length of model timesteps, $S \in \{20, 30, 60\}$. For a fixed NFE budget of 300, these corresponded to 15, 10 and 5 epochs for zero-shot fine-tuning respectively. Fig. 5a shows the final estimates, while Fig. 5b and Fig. 5c depict the corresponding loss and PSNR curves for each combination (Further quantitative results are in SuppMat). Notably, all the estimates are similar, though sharpness improves slightly as S increases. However, the trade-off for choosing a high S is the low number of epochs. Especially for cases, where the measurement system or noise level changes, this makes fine-tuning susceptible to initialization of the hyperparameters as it is more difficult to converge to a good solution in ~ 5 epochs. Thus, for improved generalizability and robustness, we opted to use $S = 30$ and 10 epochs for our database testing.


Our second ablation study analyzed the performance of ZAPS with respect to other state-of-the-art methods when all methods used the same NFE. We

Table 3: Computational costs of methods in terms of NFEs and wall-clock time (WCT)


	DPS [6]	MCG [7]	IIGDM [31]	DDRM [21]	Score-SDE [34]	ZAPS
Total NFEs	1000	1000	100	20	1000	300
WCT (s)	47.25	48.83	4.53	2.12	23.47	14.71



(a) Reconstructions using ZAPS for super-resolution ($\times 4$) task with different total timesteps-epochs combinations for the same NFE=300.



(b) Loss graphs for each combination.



(c) PSNR graphs for each combination.

Fig. 5: Study on different epochs and sampling steps combinations with fixed NFE. Results show similar quality for combinations with lower timestep approaches starting from higher loss/lower PSNR but converging to similar values.

investigated total NFEs of 100, 300, and 500 to demonstrate the robustness of our approach, given its adaptable parameters, as previously discussed. For 100 NFEs, we applied 20 steps (schedule = “10, 7, 3”) with 5 epochs, whereas for 300 and 500 NFEs, we applied 30 steps (schedule = “15, 10, 5”) and 50 steps (schedule = “30, 15, 5”), respectively, for 10 epochs. Additionally, we also implemented ZAPS with uniformly spaced noise schedules to highlight the benefits of the proposed irregular noise schedules. As seen in Tabs. 4 and 5, ZAPS with irregular noise schedules outperforms the state-of-the-art methods for NFE budgets of 100, 300 and 500 in super-resolution and random inpainting tasks. We note that we could not perform this test for deblurring experiments as IIGDM could not be implemented reliably across the database, as previously mentioned. We also note that the difference between irregular and uniform noise schedules for ZAPS is

Table 4: Quantitative results for super-resolution ($\times 4, \sigma = 0.05$) on FFHQ dataset using the same NFE for each method. Best: **bold**, second-best: underlined.

Method	NFE=100			NFE=300			NFE=500		
	LPIPS↓	SSIM↑	PSNR↑	LPIPS↓	SSIM↑	PSNR↑	LPIPS↓	SSIM↑	PSNR↑
DPS [6]	0.344	0.478	16.96	0.257	0.577	20.01	0.218	0.623	21.52
IIGDM [31]	0.131	0.760	24.48	<u>0.117</u>	<u>0.758</u>	24.80	0.123	<u>0.762</u>	24.25
ZAPS (Uniform)	<u>0.108</u>	<u>0.749</u>	<u>25.92</u>	0.119	0.729	<u>26.29</u>	<u>0.115</u>	0.756	<u>25.63</u>
ZAPS (Irregular)	0.106	0.741	26.08	0.104	0.768	26.63	0.095	0.770	26.26

Table 5: Quantitative results for random inpainting (70%, $\sigma = 0.05$) on FFHQ dataset using the same NFE for each method. Best: **bold**, second-best: underlined.

Method	NFE=100			NFE=300			NFE=500		
	LPIPS↓	SSIM↑	PSNR↑	LPIPS↓	SSIM↑	PSNR↑	LPIPS↓	SSIM↑	PSNR↑
DPS [6]	0.268	0.593	20.01	0.189	0.704	23.74	0.152	0.754	25.59
IGDM [31]	<u>0.086</u>	0.842	<u>26.62</u>	<u>0.080</u>	0.849	25.06	0.082	0.845	24.94
ZAPS (Uniform)	0.122	0.780	26.20	0.127	0.773	<u>25.87</u>	<u>0.080</u>	0.791	<u>26.94</u>
ZAPS (Irregular)	0.085	<u>0.794</u>	27.03	0.078	<u>0.813</u>	27.79	0.071	<u>0.818</u>	28.11

less pronounced for 100 NFEs, but the advantage of irregular schedules becomes apparent for 300 and 500 NFEs.

The final ablation study, exploring the benefits of using distinct weights ζ_t for each timestep versus a shared weight ζ for every step, is provided in SuppMat.

4.4 Limitations

The loss function we use, $\mathcal{L}(\mathbf{y}, \mathbf{x}_0) = \|\mathbf{y} - \mathbf{A}\mathbf{x}_0\|_2^2$, resembles a deep image prior-like loss [38]. However, note that there is a subtle difference in our context, where it corresponds to the log-likelihood of $p(\mathbf{y}|\mathbf{x}_0)$, which is different than the (approximate) log-likelihood guidance term $p(\mathbf{y}|\mathbf{x}_t)$ used at each time-step. This allows for more robustness to overfitting that is typically observed in DIP-type methods. Further overfitting avoidance measures can be taken by data-splitting [3, 23, 26, 41, 42], though this was not necessary for the small number of epochs used for fine-tuning. Additionally, while our approximation in Eq. (22) produces competitive results, it is important to keep in mind that wavelets may not fully decorrelate the observed Fisher information matrix. Finally, we note that while we chose DPS as a baseline for its versatility in inverse problem tasks, the adaptive weighting strategy in ZAPS, as well as our Hessian approximation, are applicable to other posterior sampling diffusion models for inverse problems.

5 Conclusion

In this work, we proposed a novel approach named zero-shot approximate posterior sampling (ZAPS), which harnesses zero-shot learning for dynamic automated hyperparameter tuning during the inference phase to enhance the reconstruction quality of solving linear noisy inverse problems using diffusion models. In particular, learning the log-likelihood weights facilitates the usage of more complex and irregular noise schedules, whose feasibility for inverse problems was shown, to the best of our knowledge, for the first time in this paper. These irregular noise schedules enabled high quality reconstructions with $20 - 50\times$ fewer timesteps. When number of epochs for fine-tuning is also considered, our approach results in a speed boost of approximately $3\times$ compared to state-of-the-art methods like DPS. Quantitative and qualitative evaluations on natural images illustrate our method’s ability to attain state-of-the-art performance across diverse inverse problem tasks.

Acknowledgements

This work was partially supported by NIH R01HL153146 and NIH R01EB032830.

References

1. Alcaraz, J.M.L., Strodthoff, N.: Diffusion-based time series imputation and forecasting with structured state space models. arXiv preprint arXiv:2208.09399 (2022)
2. Baranchuk, D., Rubachev, I., Voynov, A., Khruikov, V., Babenko, A.: Label-efficient semantic segmentation with diffusion models. International Conference on Learning Representations (2021)
3. Batson, J., Royer, L.: Noise2self: Blind denoising by self-supervision. In: International Conference on Machine Learning. pp. 524–533. PMLR (2019)
4. Chan, S.H., Wang, X., Elgandy, O.A.: Plug-and-play admm for image restoration: Fixed-point convergence and applications. IEEE Transactions on Computational Imaging **3**(1), 84–98 (2016)
5. Choi, J., Kim, S., Jeong, Y., Gwon, Y., Yoon, S.: Ilvr: Conditioning method for denoising diffusion probabilistic models. in 2021 IEEE. In: CVF international conference on computer vision (ICCV). pp. 14347–14356 (2021)
6. Chung, H., Kim, J., McCann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. International Conference on Learning Representations (2023)
7. Chung, H., Sim, B., Ryu, D., Ye, J.C.: Improving diffusion models for inverse problems using manifold constraints. Advances in Neural Information Processing Systems (2022)
8. Chung, H., Sim, B., Ye, J.C.: Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022)
9. Cohen, R., Blau, Y., Freedman, D., Rivlin, E.: It has potential: Gradient-driven denoisers for convergent solutions to inverse problems. Advances in Neural Information Processing Systems **34**, 18152–18164 (2021)
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
11. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in neural information processing systems **34**, 8780–8794 (2021)
12. Ghael, S., Sayeed, A.M., Baraniuk, R.G.: Improved wavelet denoising via empirical wiener filtering. In: SPIE Technical Conference on Wavelet Applications in Signal Processing (1997)
13. Graikos, A., Malkin, N., Jovic, N., Samaras, D.: Diffusion models as plug-and-play priors. Advances in Neural Information Processing Systems **35**, 14715–14728 (2022)
14. Gregor, K., LeCun, Y.: Learning fast approximations of sparse coding. In: Proceedings of the 27th international conference on international conference on machine learning. pp. 399–406 (2010)
15. Hammernik, K., Küstner, T., Yaman, B., Huang, Z., Rueckert, D., Knoll, F., Akçakaya, M.: Physics-driven deep learning for computational magnetic resonance imaging. IEEE Sig Proc Mag **40**, 98–114 (2023)

16. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
17. Hoogeboom, E., Nielsen, D., Jaini, P., Forré, P., Welling, M.: Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in Neural Information Processing Systems* **34**, 12454–12465 (2021)
18. Kadkhodaie, Z., Simoncelli, E.: Stochastic solutions for linear inverse problems using the prior implicit in a denoiser. *Advances in Neural Information Processing Systems* **34**, 13242–13254 (2021)
19. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems* **35**, 26565–26577 (2022)
20. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 4396–4405 (2019)
21. Kawar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. In: *Advances in Neural Information Processing Systems* (2022)
22. Knoll, F., Hammernik, K., Zhang, C., Moeller, S., Pock, T., Sodickson, D.K., Akçakaya, M.: Deep learning methods for parallel magnetic resonance imaging reconstruction. *IEEE Sig Proc Mag* **37**, 128–140 (2020)
23. Krull, A., Buchholz, T.O., Jug, F.: Noise2void-learning denoising from single noisy images. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 2129–2137 (2019)
24. Laumont, R., Bortoli, V.D., Almansa, A., Delon, J., Durmus, A., Pereyra, M.: Bayesian imaging using plug & play priors: when langevin meets tweedie. *SIAM Journal on Imaging Sciences* **15**(2), 701–737 (2022)
25. Mardani, M., Song, J., Kautz, J., Vahdat, A.: A variational perspective on solving inverse problems with diffusion models. *arXiv preprint arXiv:2305.04391* (2023)
26. Moran, N., Schmidt, D., Zhong, Y., Coady, P.: Noisier2noise: Learning to denoise from unpaired noisy data. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12064–12072 (2020)
27. Qu, Y., Zheng, N., Li, C.: Using wavelet transform to estimate the eigenfunctions of karhunen-loeve expansion. In: *Wavelet Analysis and Its Applications, and Active Media Technology*, pp. 39–44. World Scientific (2004)
28. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *International conference on machine learning*. pp. 2256–2265. PMLR (2015)
29. Song, B., Kwon, S.M., Zhang, Z., Hu, X., Qu, Q., Shen, L.: Solving inverse problems with latent diffusion models via hard data consistency. *arXiv preprint arXiv:2307.08123* (2023)
30. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *International Conference on Learning Representations* (2020)
31. Song, J., Vahdat, A., Mardani, M., Kautz, J.: Pseudoinverse-guided diffusion models for inverse problems. In: *International Conference on Learning Representations* (2022)
32. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems* **32** (2019)
33. Song, Y., Shen, L., Xing, L., Ermon, S.: Solving inverse problems in medical imaging with score-based generative models. *arXiv preprint arXiv:2111.08005* (2021)
34. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations* (2020)

35. Sun, Y., Wang, X., Liu, Z., Miller, J., Efros, A., Hardt, M.: Test-time training with self-supervision for generalization under distribution shifts. In: International conference on machine learning. pp. 9229–9248. PMLR (2020)
36. Taam, W., Yandell, B.S.: Approximate Diagonalization of Spatial Covariance. University of Wisconsin, Department of Statistics (1987)
37. Tumanyan, N., Geyer, M., Bagon, S., Dekel, T.: Plug-and-play diffusion features for text-driven image-to-image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1921–1930 (2023)
38. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9446–9454 (2018)
39. Vincent, P.: A connection between score matching and denoising autoencoders. *Neural computation* **23**(7), 1661–1674 (2011)
40. Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. The Eleventh International Conference on Learning Representations (2023)
41. Yaman, B., Hosseini, S.A.H., Moeller, S., Ellermann, J., Uğurbil, K., Akçakaya, M.: Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data. *Magn Reson Med* **84**(6), 3172–3191 (Dec 2020)
42. Yaman, B., Hosseini, S.A.H., Akçakaya, M.: Zero-shot self-supervised learning for MRI reconstruction. *Proc ICLR* (2021)
43. Yang, L., Ding, S., Cai, Y., Yu, J., Wang, J., Shi, Y.: Guidance with spherical gaussian constraint for conditional diffusion. In: International Conference on Machine Learning (2024)
44. Zhu, Y., Zhang, K., Liang, J., Cao, J., Wen, B., Timofte, R., Gool, L.V.: Denoising diffusion models for plug-and-play image restoration. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops (NTIRE) (2023)