

Local Epistemic Uncertainty Guided Active Sampling for Plug-and-play Diffusive Image Restoration

Anonymous Author(s)
Submission Id: 7960

Abstract

Diffusion models have demonstrated remarkable effectiveness in image restoration tasks. However, when guiding image reconstruction, existing Diffusion Model-based Image Restoration (DMIR) methods typically rely on fixed data constraints and uniform step sizes, thereby overlooking the dynamic nature of the generative process. Such rigid designs render the models vulnerable to spatially non-uniform degradations, thus resulting in structural distortions and loss of fine details. Meanwhile, uniform step sizes introduce computational redundancy, whereas naïve step reduction strategies tend to accumulate approximation errors. To address these limitations, we propose a Local Epistemic Uncertainty Guided Active Sampling framework (LEADer). In the spatial domain, LEADer leverages pixel-wise uncertainty to dynamically modulate the prior strength within the null space, which effectively balances detail preservation and artifact suppression. In the temporal domain, it quantifies sampling stability via the uncertainty trace to enable adaptive trajectory pruning, thereby accelerating convergence. Theoretical proofs demonstrate that our framework achieves strict data consistency, while the trajectory pruning strategy admits a deterministic error bound, thereby guaranteeing stable convergence under skip sampling. Notably, our plug-and-play method can be seamlessly integrated into various DMIR baselines. Extensive experiments show that LEADer improves the performance of multiple state-of-the-art DMIR methods, while significantly reducing sampling time with negligible memory overhead.

CCS Concepts

• **Computing methodologies** → **Computer vision**; *Machine learning*.

Keywords

Diffusion model, Plug-and-play, Image restoration, Local epistemic uncertainty

1 Introduction

Image restoration [9, 11, 15, 20, 22] aims to reconstruct clear, natural images from degraded observations corrupted by noise, blur, or low resolution, serving as a fundamental problem in computer vision and medical imaging [27, 40, 42]. Recently, deep learning-based methods have achieved remarkable performance on various specific image restoration tasks [4, 11, 22, 48, 52]. However, these methods, optimized for specific degradation patterns, typically rely on massive paired training data. Moreover, this task-specific paradigm often suffers from limited generalization capabilities [51].

Recently, diffusion models [8, 17, 31, 34, 36, 46] have demonstrated powerful capabilities in modeling complex data distributions and have been successfully applied to diverse image generation

tasks [10, 29]. Consequently, researchers have begun exploring pre-trained unconditional diffusion models as universal priors to solve image restoration inverse problems [13]. Diffusion Model-based Image Restoration (DMIR) methods typically adopt a posterior sampling paradigm: they directly utilize pre-trained diffusion models to fit the prior distribution of clear natural images, incorporating likelihood guidance [2, 5, 7, 26, 35, 39, 44, 45] or data consistency projection [6, 14, 19, 25, 38, 41, 43, 47, 50, 54] during the reverse sampling process. This enables a single pre-trained model to adapt to diverse degradation scenarios, which effectively enhances the generalization capability of restoration methods.

However, during reverse sampling, existing DMIR methods typically employ fixed data constraint weights and uniform time steps. This setting overlooks the dynamic nature of image generation and introduces two main challenges. Spatially, since local degradations in real images are often uneven, globally fixed weights struggle to accommodate varying regional restoration needs, thus leading to compromised details or inaccurate local structures. Temporally, fixed time steps cause additional computational cost during the stationary phase of sampling, whereas directly reducing steps introduces additional discretization errors that compromise overall restoration quality [18, 24].

In response to these challenges, we propose an active diffusion sampling framework guided by local epistemic uncertainty, termed LEADer. Specifically, to ensure structural consistency, we propose Uncertainty-Calibrated Prior Modulation (UCPM), which leverages pixel-level uncertainty to dynamically regulate the strength of prior within the null space, thereby adaptively balancing detail preservation and artifact suppression. Furthermore, to improve inference efficiency, we propose State-Aware Trajectory Pruning (SATP). This technique evaluates the stability of the sampling process via the trace of uncertainty and executes step-skipping within a controlled error tolerance, which effectively reduces computational redundancy and accelerates convergence, as illustrated in Figure 1. In addition, theoretical analysis demonstrates that our method guarantees data fidelity while providing a controlled error bound for the skip sampling process. The main contributions of this paper are summarized as follows:

- We propose a novel perspective on the diffusion reverse process based on local epistemic uncertainty, which reveals the spatiotemporal limitations of the sampling strategies in existing DMIR methods.
- We propose an active diffusion sampling framework that proactively resolves local distortion and sampling redundancy via uncertainty-calibrated prior modulation and state-aware trajectory pruning.
- Extensive experimental results demonstrate that our proposed plug-and-play framework effectively enhances image restoration quality and performance while improving inference efficiency.

2 Background

DMIR methods have emerged as the mainstream for zero-shot solvers in image inverse problems. The core challenge lies in injecting specific measurement constraints into the reverse denoising process. Depending on the intervention mechanisms, existing methods are primarily categorized into the following two classes:

Gradient-based Optimization Guidance: Rooted in Bayesian posterior sampling theory [5, 26, 35], these methods approximate the measurement likelihood gradient using the clean image estimate at the current time step to guide reverse sampling. Notably, DPS [5] directly enforces data consistency using the gradient of the l_2 -norm of the measurement error. Subsequent works have focused on improving guidance accuracy and stability. For example, MCG [7] incorporates manifold projection constraints, while IIGDM [33] enhances robustness via pseudo-inverse matrices and Jacobian computations. ZAPS [1] further proposes an alternative guidance scheme to improve performance over DPS [5]. More recent studies investigate the numerical instability and bias arising during the guidance process. For instance, SITCOM [2] and SPGD [45] introduce explicit gradient control mechanisms to stabilize optimization. In parallel, DCDP [21] and DAPS [49] refine the sampling trajectory by decoupling the purification process to mitigate prior bias, while DPPS [44] employs proximal operators to handle non-smooth constraints, thereby improving flexibility under complex degradations.

Null-Space Orthogonal Projection: In contrast to gradient-based guidance, these methods decompose the image space into range and null spaces via orthogonal projection, and iteratively refine the null-space components during reverse diffusion. Representative approaches such as DDNM [41] and DDRM [19] achieve strict data consistency through pseudo-inverse projection and singular value decomposition (SVD), respectively. Building upon this foundation, subsequent works extend the framework to improve flexibility and guidance quality. DiffPIR [54] integrates Half-Quadratic Splitting (HQS) to reformulate measurement constraints as proximal mappings, while DDPG [14] introduces a preconditioned guidance trajectory that evolves with sampling steps, which effectively bridges back-projection and least-squares gradient formulations. More recent studies focus on enhancing global consistency and structural regularity. For example, ProjDiff [50] constructs a bivariate constrained optimization framework to better exploit denoising priors, EquS [43] utilizes equivariant sampling to constrain manifold regularity, and PIRP [47] incorporates parameterized gradient priors to elevate overall perception and fidelity.

3 Preliminary

Problem Formulation: The image restoration task aims to recover a target clean image $\mathbf{x}_0 \in \mathbb{R}^N$ from a degraded observation $\mathbf{y} \in \mathbb{R}^M$. This physical degradation process is typically modeled as a linear inverse problem:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}, \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^M$ is the degraded image, and $\mathbf{x}_0 \in \mathbb{R}^N$ represents the high-quality image to be recovered. $\mathbf{A} \in \mathbb{R}^{M \times N}$ is a known linear operator, such as a bicubic downsampler in image super-resolution. Finally, $\mathbf{n} \sim \mathcal{N}(0, \sigma_y^2 \mathbf{I})$ is additive Gaussian noise [19], which represents uncertainty or disturbance in the observation \mathbf{y} .

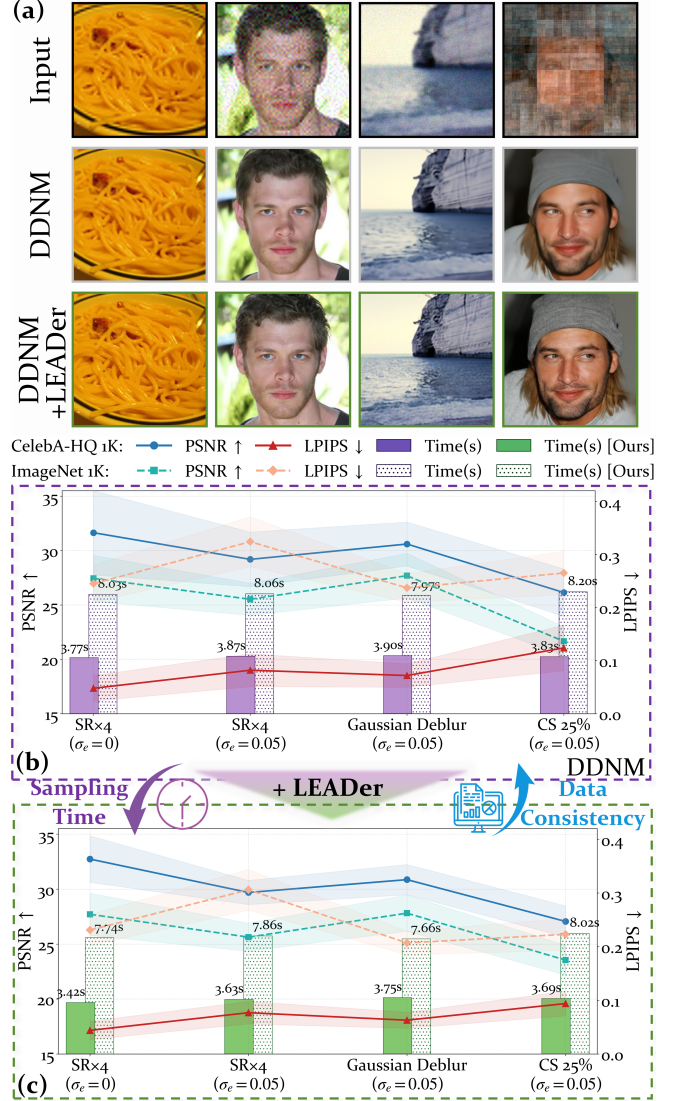


Figure 1: (a) Qualitative comparisons of the proposed method with plug-and-play adaptation across different tasks. (b) and (c) Quantitative comparisons under the same settings. The results demonstrate that our method facilitates restoration accuracy and efficiency.

Within the zero-shot diffusion framework, restoration is achieved without retraining on paired degraded data. Instead, a pre-trained unconditional diffusion model serves as a generative prior. At any time step $t \in (0, T)$ during reverse sampling, given the current noisy variable \mathbf{x}_t , the diffusion model predicts the noise component via a denoising neural network $\epsilon_\theta(\mathbf{x}_t, t)$. Based on Tweedie's formula [12, 37], the corresponding estimate of the clean image is:

$$\mathbf{x}_{0|t} = \frac{\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t)}{\sqrt{\bar{\alpha}_t}}, \quad (2)$$

where $\bar{\alpha}_t$ is a predefined noise schedule parameter.

To enforce data consistency with the physical observation, the Denoising Diffusion Null-Space Model [41] introduces an orthogonal decomposition of the image space into range and null spaces:

$$\hat{\mathbf{x}}_{0|t} = \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) \mathbf{x}_{0|t}, \quad (3)$$

where \mathbf{A}^\dagger is the pseudo-inverse of the degradation matrix. $\mathbf{A}^\dagger \mathbf{y}$ enforces data fidelity in the range space, while $(\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) \mathbf{x}_{0|t}$ projects $\mathbf{x}_{0|t}$ onto the null space for prior-guided refinement.

Diffusion Models: Denoising Diffusion Probabilistic Models (DDPM) [17] learn data generation by reversing a diffusion process that progressively corrupts data through noise addition. This framework consists of a forward process and a reverse process.

The forward process is a fixed Markov chain that gradually perturbs the data over T time steps. Let β_t denote the noise schedule. At each step t , the noisy state \mathbf{x}_t is given by:

$$\mathbf{x}_t = \sqrt{1 - \beta_t} \mathbf{x}_{t-1} + \sqrt{\beta_t} \boldsymbol{\epsilon}_{t-1}, \quad (4)$$

where $\boldsymbol{\epsilon}_{t-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Let the cumulative noise up to time step t be $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, then we have:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}. \quad (5)$$

The reverse process aims to reconstruct the original clean image \mathbf{x}_0 from the noisy state \mathbf{x}_t . Mathematically, this is described by:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) \right) + \sigma_t \boldsymbol{\epsilon}_t, \quad (6)$$

where $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)$ denotes the noise predicted by a neural network parameterized by θ , and $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is the Gaussian noise.

Despite their strong generative capability, DDPM suffers from high computational cost due to the large number of sampling steps. To address this limitation, [32] introduces Denoising Diffusion Implicit Models (DDIM), which relax the Markovian assumption, and enable more efficient sampling with fewer steps. Based on the estimate $\mathbf{x}_{0|t}$ in Eq. (2), the DDIM update rule is:

$$\mathbf{x}_{t-1} = \sqrt{\bar{\alpha}_{t-1}} \mathbf{x}_{0|t} + \sqrt{1 - \bar{\alpha}_{t-1} - \sigma_t^2} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) + \sigma_t \boldsymbol{\epsilon}_t. \quad (7)$$

4 Local Epistemic Uncertainty Guided Active Sampling

4.1 Motivation

Existing zero-shot image restoration solvers typically employ globally fixed prior weights and uniform time steps when intervening in the reverse diffusion process. Such a rigid design implicitly treats sampling as a static procedure and overlooks the inherently dynamic characteristics of diffusion models across different spatial regions and temporal evolution stages. This limitation gives rise to two key issues:

Perception-Distortion Trade-off: The level of degradation and the difficulty of restoration varies significantly across different image regions [3]. Applying globally uniform constraint weights fails to account for this heterogeneity. Specifically, in regions where the model can already recover fine details, overly strong priors tend to over-smooth structures. On the other hand, in severely degraded areas where prior-driven inference is less reliable, insufficient constraints can lead to artifacts that deviate from the true underlying structures [16, 28].

Sampling Redundancy: From the perspective of ordinary differential equation (ODE), the reverse diffusion trajectory evolves with non-uniform dynamics. However, existing methods typically adopt fixed, uniform step sizes. This static temporal scheduling ignores the varying demands of different sampling stages. Once the global structure is established, continuing dense step-by-step updates introduces substantial computational redundancy; whereas naively increasing step sizes can accumulate errors during detail reconstruction, which ultimately degrades the quality [18].

To address these issues, we argue that the diffusion sampling steps should be state-aware. Specifically, we quantify local epistemic uncertainty at each spatial location and sampling step, then we use it as a unified modulation signal. In the spatial dimension, this uncertainty facilitates pixel-level on-demand regularization modulation to balance detail preservation and artifact suppression (see Section 4.3). In the temporal dimension, it serves as an error control metric to guide adaptive step sizing (see Section 4.4). The overall architecture of our method is illustrated in Figure 2.

4.2 Local Epistemic Uncertainty Quantification

At any time step t in the reverse dynamics, the noisy observation \mathbf{x}_t embodies two fundamentally different types of uncertainty: the inherent aleatoric uncertainty introduced by the forward Gaussian diffusion process, and the epistemic uncertainty arising from missing information, particularly within the null space. In this paper, we focus on quantifying the latter from an information geometry perspective.

Based on the marginal distribution $q(\mathbf{x}_t | \mathbf{x}_0)$ of the forward noise-adding process, the Observed Fisher Information Matrix (FIM) $\mathcal{F}_t(\mathbf{x}_t) \in \mathbb{R}^{N \times N}$ of the pre-trained diffusion model at the current state \mathbf{x}_t is defined as:

$$\mathcal{F}_t(\mathbf{x}_t) = -\nabla_{\mathbf{x}_t}^2 \log p_t(\mathbf{x}_t), \quad (8)$$

where $\mathcal{F}_t(\mathbf{x}_t)$ represents the model's information sensitivity and structural certainty regarding the local features at this state. A higher value indicates that the model has stronger prior confidence in high-frequency details, whereas a lower value implies higher generative ambiguity in that region. As detailed in Section 3, the posterior expectation estimate $\mathbf{x}_{0|t}$ of the target clean image \mathbf{x}_0 is:

$$\mathbf{x}_{0|t} = \mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t] = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(\mathbf{x}_t + (1 - \bar{\alpha}_t) \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \right). \quad (9)$$

To quantify the prediction variance of the model given \mathbf{x}_t , we perform a second-order Taylor expansion on the posterior distribution to derive the analytical relationship between the posterior covariance matrix Σ_t and FIM:

$$\Sigma_t = \text{Cov}(\mathbf{x}_0 | \mathbf{x}_t) = \frac{1 - \bar{\alpha}_t}{\bar{\alpha}_t} \left[\mathbf{I} - (1 - \bar{\alpha}_t) \mathcal{F}_t(\mathbf{x}_t) \right]. \quad (10)$$

This decomposition provides a clear interpretation. The leading term $\frac{1 - \bar{\alpha}_t}{\bar{\alpha}_t} \mathbf{I}$ represents the inherent upper bound of the aleatoric uncertainty at time step t , while the second term $(1 - \bar{\alpha}_t) \mathcal{F}_t(\mathbf{x}_t)$ reflects the Fisher information gain provided by the pre-trained model via the score function. Consequently, the diagonal matrix Σ_t serve as an effective measure of local epistemic uncertainty. Specifically, when $\mathcal{F}_t(\mathbf{x}_t)$ is large, the model exhibits high confidence in the reconstructed structures, thus leading to an epistemic uncertainty

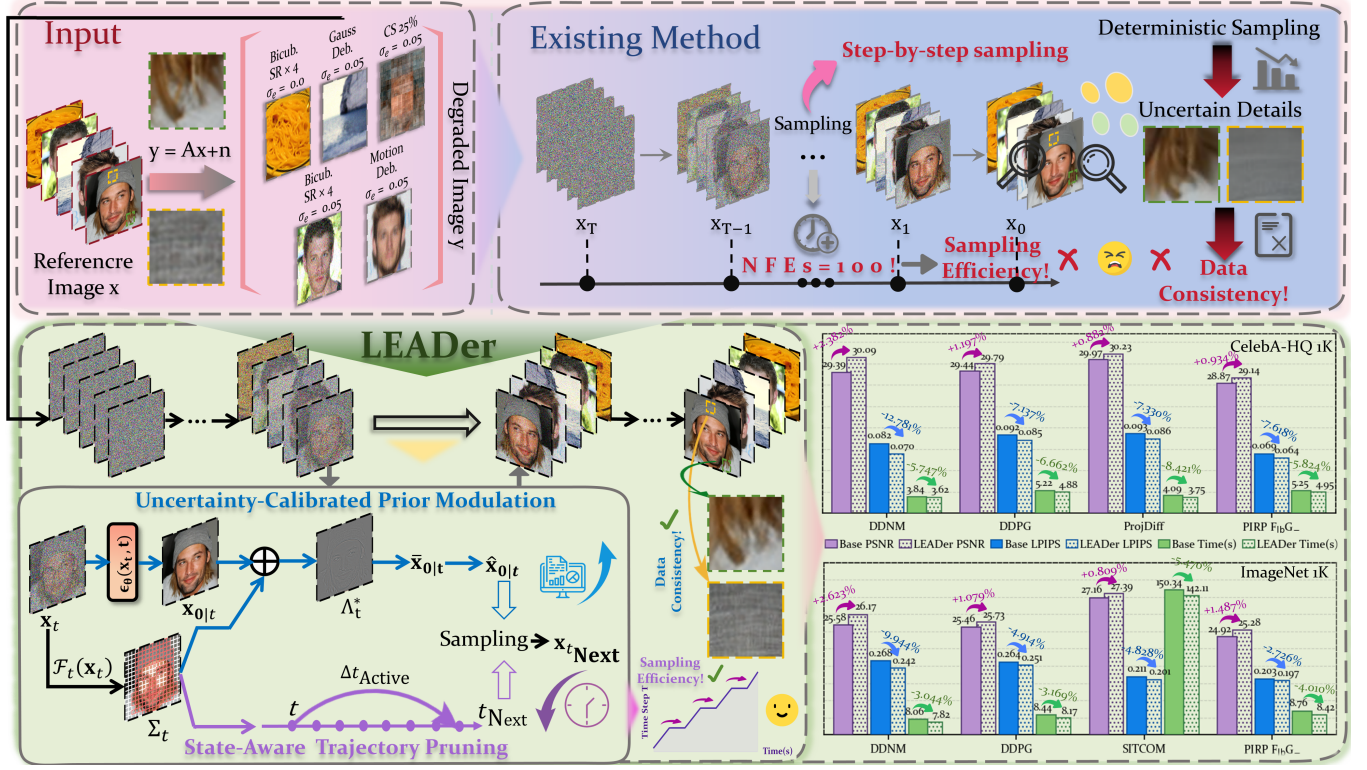


Figure 2: Architecture of LEADER. Existing DMIR methods rely on fixed step-by-step sampling with uniform constraints, leading to computational redundancy and detail distortion in uncertain regions. To address this, LEADER quantifies local epistemic uncertainty Σ_t to guide reverse sampling. Spatially, UCPM utilizes Σ_t to adaptively modulate prior constraints for strict data consistency. Temporally, SATP computes an adaptive step size Δt_{Active} based on the uncertainty trace to skip redundant iterations. As a result, LEADER improves the restoration quality of various baselines while reducing sampling time.

of $\Sigma_t \rightarrow 0$. Conversely, in regions with limited observational constraints, particularly within the null-space, the Fisher information diminishes, thus resulting in a higher Σ_t .

4.3 Uncertainty-Calibrated Prior Modulation

To address the limitations of globally uniform weighting, we propose Uncertainty-Calibrated Prior Modulation (UCPM). By leveraging the extracted local epistemic uncertainty Σ_t as a spatial guidance signal, UCPM dynamically modulates the strength of physical priors $\Phi(\mathbf{x})$ (e.g., gradient and structural priors) at the pixel level. Specifically, we cast the problem as a Maximum A Posteriori (MAP) optimization objective within the local state space:

$$\mathcal{J}(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}_{0|t})^T \Sigma_t^{-1}(\mathbf{x} - \mathbf{x}_{0|t}) + \Phi(\mathbf{x}). \quad (11)$$

Setting the first-order optimality condition $\nabla \mathcal{J}(\mathbf{x}) = \mathbf{0}$ yields:

$$\Sigma_t^{-1}(\mathbf{x} - \mathbf{x}_{0|t}) + \nabla \Phi(\mathbf{x}) = \mathbf{0}. \quad (12)$$

To obtain a tractable solution, we linearize the non-linear prior term via a first-order Taylor expansion around $\mathbf{x}_{0|t}$:

$$\nabla \Phi(\mathbf{x}) \approx \nabla \Phi(\mathbf{x}_{0|t}) + \mathbf{H}_\Phi(\mathbf{x}_{0|t})(\mathbf{x} - \mathbf{x}_{0|t}), \quad (13)$$

where $\mathbf{H}_\Phi(\mathbf{x}_{0|t}) \in \mathbb{R}^{N \times N}$ is the Hessian matrix of the prior functional. Substituting this into the optimality condition, we derive

the closed-form solution for the optimal image estimate $\bar{\mathbf{x}}_{0|t}$:

$$\bar{\mathbf{x}}_{0|t} = \mathbf{x}_{0|t} - \Lambda_t^* \nabla \Phi(\mathbf{x}_{0|t}), \quad (14)$$

where the spatially heterogeneous modulation matrix Λ_t^* is defined as:

$$\Lambda_t^* = [\Sigma_t^{-1} + \mathbf{H}_\Phi(\mathbf{x}_{0|t})]^{-1} \equiv \Sigma_t (\mathbf{I} + \mathbf{H}_\Phi(\mathbf{x}_{0|t}) \Sigma_t)^{-1}. \quad (15)$$

The matrix Λ_t^* acts as an adaptive preconditioner, with its modulation strength governed by Σ_t . When $\Sigma_t \rightarrow 0$, we have $\Lambda_t^* \rightarrow \mathbf{0}$, which indicates that prior intervention is attenuated to preserve high-frequency details generated by the diffusion model. In contrast, a larger value of Σ_t leads to stronger modulation, which enforces more aggressive regularization to suppress artifacts and structural deviations.

4.4 State-Aware Trajectory Pruning

To alleviate the computational redundancy caused by fixed time steps during the stable evolution phase, LEADER introduces State-Aware Trajectory Pruning (SATP), which incorporates epistemic uncertainty into the temporal dimension to enable adaptive sampling acceleration. In the reverse sampling of diffusion models, a single-step state transition can be approximated via Taylor expansion. Specifically, the state transition equation from time step t to

$t - \Delta t$ can be evaluated as:

$$\mathbf{x}_{t-\Delta t} = \mathbf{x}_t - \Delta t \frac{d\mathbf{x}_t}{dt} + O\left((\Delta t)^2 \left\| \frac{d^2\mathbf{x}_t}{dt^2} \right\|\right). \quad (16)$$

The higher-order residual term characterizes the local evolution error, whose magnitude reflects the rate of change in the image state. We relate this quantity to the epistemic uncertainty Σ_t derived in Section 4.2. Specifically, the norm of the second-order derivative is proportional to the trace of Σ_t :

$$\left\| \frac{d^2\mathbf{x}_t}{dt^2} \right\| \propto \text{Tr}(\mathbf{I} + (1 - \bar{\alpha}_t) \nabla_{\mathbf{x}_t}^2 \log p_t(\mathbf{x}_t)) \propto \text{Tr}(\Sigma_t) \triangleq U_t, \quad (17)$$

where $U_t = \frac{1}{N} \text{Tr}(\Sigma_t)$ denotes the global epistemic uncertainty at time step t . To balance reconstruction fidelity and computational efficiency, we introduce a constant information loss budget B that constrains the permissible local truncation error. For each step. Specifically, the error induced by a step of size Δt is bounded as:

$$\delta(\Delta t) \approx (\Delta t)^\rho \cdot U_t \leq B, \quad (18)$$

where $\rho \geq 1$ is the local error order of the sampler. In our implementation, we use the DDIM solver and set $\rho = 1$. By solving this error constraint, the maximum admissible step size Δt_{Active} can be derived, and the next evolution state node t_{Next} can be updated accordingly:

$$\Delta t_{\text{Active}} = \max \left\{ 1, \left\lfloor \eta \left(\frac{B}{U_t + \epsilon} \right)^{\frac{1}{\rho}} \right\rfloor \right\}, \quad (19)$$

$$t_{\text{Next}} = \max\{0, t - \Delta t_{\text{Active}}\},$$

where η is a scaling factor, and ϵ is a constant that ensures numerical stability. This formulation enables a principled trade-off between efficiency and accuracy. When the model is confident in the current estimate, it indicates that the image is in a stable restoration phase. At this point, the induced local truncation error remains small, thereby permitting larger step sizes and the skipping of redundant iterations. Conversely, in high-uncertainty regions associated with complex structural reconstruction, the step size is reduced to maintain reconstruction accuracy.

4.5 More Analysis

Algorithm Summary: In summary, the spatially optimal modulated prior estimate is computed via Eq.(14). To strictly ensure physical measurement consistency, we project it onto the null-space and fuse it with the range-space observation \mathbf{y} , which yields the clean image estimate:

$$\hat{\mathbf{x}}_{0|t} = \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) \bar{\mathbf{x}}_{0|t}. \quad (20)$$

Finally, following the non-Markovian formulation of DDIM, the next state $\mathbf{x}_{t_{\text{Next}}}$ can be derived using $\hat{\mathbf{x}}_{0|t}$ and the adaptively selected time step t_{Next} :

$$\mathbf{x}_{t_{\text{Next}}} = \sqrt{\bar{\alpha}_{t_{\text{Next}}}} \hat{\mathbf{x}}_{0|t} + \sqrt{1 - \bar{\alpha}_{t_{\text{Next}}} - \sigma_{t_{\text{Next}}}^2} \epsilon_\theta(\mathbf{x}_t, t) + \sigma_{t_{\text{Next}}} \epsilon_t, \quad (21)$$

where $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is the sampling noise. The overall LEADER algorithm is summarized in Algorithm 1.

Plug-and-Play Analysis: Notably, the LEADER framework can be seamlessly integrated into existing baseline sampling methods as a plug-and-play module. The computations of UCPM and SATP rely entirely on the posterior statistics ($\mathbf{x}_{0|t}$ and Σ_t) of the diffusion

Algorithm 1 Local Epistemic Uncertainty Guided Active Sampling

Require: Noise estimator $\epsilon_\theta(\cdot, t)$, T , \mathbf{y} , \mathbf{A} , Φ , B , η , ρ , ϵ

```

1: Initialize  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
2: while  $t > 0$  do
3:   // Pre-computation
4:    $\mathbf{x}_{0|t} = \frac{1}{\sqrt{\bar{\alpha}_t}} (\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t))$ 
5:    $\mathcal{F}_t(\mathbf{x}_t) = -\nabla_{\mathbf{x}_t}^2 \log p_t(\mathbf{x}_t)$ 
6:    $\Sigma_t = \frac{1 - \bar{\alpha}_t}{\bar{\alpha}_t} [\mathbf{I} - (1 - \bar{\alpha}_t) \mathcal{F}_t(\mathbf{x}_t)]$ 
7:   // Uncertainty-Calibrated Prior Modulation (UCPM)
8:    $\Lambda_t^* = \Sigma_t (\mathbf{I} + \mathbf{H}_\Phi(\mathbf{x}_{0|t}) \Sigma_t)^{-1}$ 
9:    $\bar{\mathbf{x}}_{0|t} = \mathbf{x}_{0|t} - \Lambda_t^* \nabla \Phi(\mathbf{x}_{0|t})$ 
10:   $\hat{\mathbf{x}}_{0|t} = \mathbf{A}^\dagger \mathbf{y} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) \bar{\mathbf{x}}_{0|t}$ 
11:  // State-Aware Trajectory Pruning (SATP)
12:   $U_t = \frac{1}{N} \text{Tr}(\Sigma_t)$ 
13:   $\Delta t_{\text{Active}} = \max \left\{ 1, \left\lfloor \eta \left( \frac{B}{U_t + \epsilon} \right)^{\frac{1}{\rho}} \right\rfloor \right\}$ 
14:   $t_{\text{Next}} = \max\{0, t - \Delta t_{\text{Active}}\}$ 
15:  // Generative Sampling Step
16:   $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
17:   $\mathbf{x}_{t_{\text{Next}}} = \sqrt{\bar{\alpha}_{t_{\text{Next}}}} \hat{\mathbf{x}}_{0|t} + \sqrt{1 - \bar{\alpha}_{t_{\text{Next}}} - \sigma_{t_{\text{Next}}}^2} \epsilon_\theta(\mathbf{x}_t, t) + \sigma_{t_{\text{Next}}} \epsilon_t$ 
18:   $t \leftarrow t_{\text{Next}}$ 
19: end while
20: return  $\hat{\mathbf{x}}$ 

```

model during the pure inference phase. This eliminates the need for any degradation-specific fine-tuning or retraining of the pre-trained denoising network ϵ_θ . This design ensures broad compatibility with existing zero-shot image restoration frameworks, facilitating its application across diverse inverse imaging tasks without modifying the underlying model parameters.

Theoretical Analysis: We now provide theoretical insights into the behavior of the UCPM and SATP mechanisms within the LEADER framework. In particular, we analyze how UCPM maintains data fidelity while introducing spatial prior modulation (Proposition 4.1), and how SATP guarantees a bounded global evolution error during adaptive temporal acceleration (Proposition 4.2). The proofs of the two Propositions are provided in the Supplementary Material.

PROPOSITION 4.1. *For a linear degradation operator \mathbf{A} and an observation \mathbf{y} , let the null-space projection be $P_N = \mathbf{I} - \mathbf{A}^\dagger \mathbf{A}$. Since the modulated estimation of UCPM is defined as $\hat{\mathbf{x}} = \mathbf{A}^\dagger \mathbf{y} + P_N \mathbf{x}_{\text{Prior}}$, under the noise-free assumption, the final generated image strictly satisfies the data fidelity condition:*

$$\|\mathbf{y} - \mathbf{A} \hat{\mathbf{x}}\|_2^2 = 0. \quad (22)$$

Since $\mathbf{A} P_N \equiv \mathbf{0}$, all prior modulation induced by the epistemic uncertainty Σ_t is strictly constrained within the null space. This implies that UCPM does not interfere with the known range-space components when suppressing artifacts or synthesizing details, thereby circumventing the common issue of data fidelity corruption caused by prior-driven updates.

PROPOSITION 4.2. *Let $\mathbf{x}^{\text{Dense}}$ denote the continuous exact solution of the ODE and $\mathbf{x}^{\text{LEADER}}$ the discrete solution accelerated by SATP. Assuming the drift field is L -Lipschitz continuous and SATP bounds the single-step truncation error within the budget B . Based on Grönwall's inequality, the global accumulated error possesses a deterministic*

Table 1: Quantitative results for five image restoration tasks on CelebA-HQ 1K (top) and ImageNet 1K (bottom). The Avg. Δ column reports the average performance improvements and time reductions achieved by our method over the baselines.

Method	Bicub. SR $\times 4$ ($\sigma_e = 0$)			Bicub. SR $\times 4$ ($\sigma_e = 0.05$)			Gaussian Deb. ($\sigma_e = 0.05$)			Motion Deb. ($\sigma_e = 0.05$)			CS 25% ($\sigma_e = 0.05$)			Avg. Δ		
	PSNR \uparrow	LPIPS \downarrow	Time \downarrow	PSNR \uparrow	LPIPS \downarrow	Time \downarrow	PSNR \uparrow	LPIPS \downarrow	Time \downarrow	PSNR \uparrow	LPIPS \downarrow	Time \downarrow	PSNR \uparrow	LPIPS \downarrow	Time \downarrow	PSNR \uparrow	LPIPS \downarrow	Time \downarrow
CelebA-HQ 1K																		
DDRM [19] [NearFPS2023]	31.64 $_{\pm 2.43}$	0.054 $_{\pm 0.021}$	3.87 s	29.26 $_{\pm 1.86}$	0.090 $_{\pm 0.028}$	4.05 s	30.53 $_{\pm 1.74}$	0.074 $_{\pm 0.024}$	4.06 s	N/A	N/A	N/A	26.07 $_{\pm 2.08}$	0.128 $_{\pm 0.037}$	3.86 s	N/A	N/A	N/A
DPS [5] [ICLR2023]	29.39 $_{\pm 3.16}$	0.065 $_{\pm 0.026}$	134.20 s	27.49 $_{\pm 2.84}$	0.086 $_{\pm 0.032}$	138.25 s	27.75 $_{\pm 2.57}$	0.084 $_{\pm 0.029}$	143.69 s	19.63 $_{\pm 3.41}$	0.227 $_{\pm 0.043}$	148.69 s	N/A	N/A	N/A	N/A	N/A	N/A
DiffPIR [54] [CVPR2023]	30.36 $_{\pm 2.27}$	0.051 $_{\pm 0.019}$	7.45 s	27.44 $_{\pm 1.93}$	0.085 $_{\pm 0.027}$	7.68 s	28.89 $_{\pm 1.78}$	0.074 $_{\pm 0.022}$	7.66 s	27.96 $_{\pm 1.61}$	0.102 $_{\pm 0.031}$	7.58 s	N/A	N/A	N/A	N/A	N/A	N/A
IDPG [14] [CVPR2024]	32.66 $_{\pm 2.18}$	0.111 $_{\pm 0.023}$	4.75 s	29.89 $_{\pm 1.64}$	0.155 $_{\pm 0.051}$	4.88 s	31.08 $_{\pm 1.52}$	0.150 $_{\pm 0.027}$	4.81 s	29.73 $_{\pm 1.73}$	0.134 $_{\pm 0.025}$	5.59 s	26.26 $_{\pm 1.95}$	0.143 $_{\pm 0.033}$	5.00 s	N/A	N/A	N/A
DDNM [41] [ICLR2023]	31.64 $_{\pm 3.87}$	0.048 $_{\pm 0.024}$	3.77 s	29.19 $_{\pm 2.51}$	0.082 $_{\pm 0.031}$	3.87 s	30.60 $_{\pm 2.02}$	0.072 $_{\pm 0.022}$	3.90 s	N/A	N/A	N/A	26.14 $_{\pm 2.18}$	0.124 $_{\pm 0.043}$	3.83 s	+2.382%	-12.781%	-5.747%
+ EquS* [43] [WACV2026]	32.44 $_{\pm 2.58}$	0.053 $_{\pm 0.019}$	3.83 s	29.20 $_{\pm 1.38}$	0.093 $_{\pm 0.027}$	3.87 s	30.83 $_{\pm 1.56}$	0.069 $_{\pm 0.021}$	3.96 s	N/A	N/A	N/A	26.31 $_{\pm 2.45}$	0.119 $_{\pm 0.035}$	4.02 s			
+ LEADer [Ours]	32.73 $_{\pm 2.12}$	0.044 $_{\pm 0.017}$	3.42 s	29.69 $_{\pm 1.10}$	0.077 $_{\pm 0.021}$	3.63 s	30.86 $_{\pm 1.38}$	0.063 $_{\pm 0.015}$	3.73 s	N/A	N/A	N/A	27.06 $_{\pm 1.42}$	0.094 $_{\pm 0.023}$	3.69 s			
DDPG [14] [CVPR2024]	31.60 $_{\pm 2.36}$	0.052 $_{\pm 0.021}$	4.96 s	29.39 $_{\pm 1.74}$	0.109 $_{\pm 0.029}$	5.09 s	30.41 $_{\pm 1.58}$	0.068 $_{\pm 0.023}$	5.11 s	29.02 $_{\pm 1.69}$	0.082 $_{\pm 0.026}$	5.72 s	26.76 $_{\pm 1.95}$	0.152 $_{\pm 0.034}$	5.24 s	+1.197%	-7.137%	-6.662%
+ EquS [43] [WACV2026]	31.73 $_{\pm 2.08}$	0.054 $_{\pm 0.019}$	4.88 s	29.52 $_{\pm 1.53}$	0.109 $_{\pm 0.027}$	5.01 s	30.47 $_{\pm 1.42}$	0.071 $_{\pm 0.021}$	5.03 s	29.17 $_{\pm 1.55}$	0.088 $_{\pm 0.024}$	5.62 s	26.82 $_{\pm 1.76}$	0.155 $_{\pm 0.032}$	5.15 s			
+ LEADer [Ours]	31.94 $_{\pm 1.79}$	0.047 $_{\pm 0.017}$	4.63 s	29.65 $_{\pm 1.31}$	0.098 $_{\pm 0.023}$	4.75 s	30.77 $_{\pm 1.24}$	0.062 $_{\pm 0.018}$	4.77 s	29.40 $_{\pm 1.37}$	0.076 $_{\pm 0.021}$	5.34 s	27.17 $_{\pm 1.58}$	0.144 $_{\pm 0.028}$	4.89 s			
ProjDiff [50] [NearFPS2024]	32.57 $_{\pm 1.88}$	0.091 $_{\pm 0.022}$	3.94 s	29.49 $_{\pm 1.42}$	0.080 $_{\pm 0.024}$	4.12 s	31.41 $_{\pm 1.35}$	0.068 $_{\pm 0.020}$	4.27 s	N/A	N/A	N/A	26.41 $_{\pm 1.57}$	0.132 $_{\pm 0.031}$	4.05 s	+0.882%	-7.330%	-8.421%
+ EquS* [43] [WACV2026]	32.64 $_{\pm 1.63}$	0.094 $_{\pm 0.020}$	3.87 s	29.53 $_{\pm 1.21}$	0.080 $_{\pm 0.021}$	4.04 s	31.45 $_{\pm 1.18}$	0.071 $_{\pm 0.018}$	4.18 s	N/A	N/A	N/A	26.45 $_{\pm 1.39}$	0.135 $_{\pm 0.028}$	3.97 s			
+ LEADer [Ours]	32.86 $_{\pm 1.34}$	0.084 $_{\pm 0.017}$	3.61 s	29.75 $_{\pm 1.05}$	0.074 $_{\pm 0.018}$	3.77 s	31.64 $_{\pm 0.98}$	0.062 $_{\pm 0.015}$	3.90 s	N/A	N/A	N/A	26.68 $_{\pm 1.16}$	0.125 $_{\pm 0.023}$	3.72 s			
PIRP F ₀ G. [47] [ORS2026]	30.80 $_{\pm 1.11}$	0.042 $_{\pm 0.018}$	5.01 s	28.55 $_{\pm 1.57}$	0.079 $_{\pm 0.024}$	5.14 s	29.90 $_{\pm 1.46}$	0.061 $_{\pm 0.020}$	5.12 s	28.50 $_{\pm 1.63}$	0.074 $_{\pm 0.022}$	5.77 s	26.59 $_{\pm 1.47}$	0.091 $_{\pm 0.027}$	5.22 s	+0.934%	-7.618%	-5.824%
+ EquS [43] [WACV2026]	30.85 $_{\pm 1.86}$	0.043 $_{\pm 0.016}$	4.94 s	28.60 $_{\pm 1.39}$	0.078 $_{\pm 0.021}$	5.06 s	29.94 $_{\pm 1.28}$	0.062 $_{\pm 0.018}$	5.04 s	28.55 $_{\pm 1.47}$	0.075 $_{\pm 0.020}$	5.67 s	26.63 $_{\pm 1.58}$	0.092 $_{\pm 0.024}$	5.14 s			
+ LEADer [Ours]	31.02 $_{\pm 1.61}$	0.039 $_{\pm 0.014}$	4.72 s	28.86 $_{\pm 1.18}$	0.072 $_{\pm 0.018}$	4.84 s	30.16 $_{\pm 1.09}$	0.055 $_{\pm 0.015}$	4.83 s	28.77 $_{\pm 1.26}$	0.069 $_{\pm 0.017}$	5.43 s	26.87 $_{\pm 1.37}$	0.086 $_{\pm 0.021}$	4.91 s			
ImageNet 1K																		
DDRM [19] [NearFPS2023]	27.38 $_{\pm 2.12}$	0.270 $_{\pm 0.045}$	8.16 s	25.54 $_{\pm 1.52}$	0.333 $_{\pm 0.047}$	8.27 s	27.71 $_{\pm 1.83}$	0.243 $_{\pm 0.032}$	8.02 s	N/A	N/A	N/A	21.58 $_{\pm 1.75}$	0.301 $_{\pm 0.038}$	8.07 s	N/A	N/A	N/A
DPS [5] [ICLR2023]	25.56 $_{\pm 1.85}$	0.236 $_{\pm 0.029}$	246.14 s	24.05 $_{\pm 1.43}$	0.271 $_{\pm 0.045}$	263.95 s	26.64 $_{\pm 1.57}$	0.240 $_{\pm 0.032}$	248.26 s	17.52 $_{\pm 1.78}$	0.468 $_{\pm 0.066}$	275.41 s	N/A	N/A	N/A	N/A	N/A	N/A
DiffPIR [54] [CVPR2023]	26.99 $_{\pm 1.92}$	0.255 $_{\pm 0.030}$	15.40 s	24.65 $_{\pm 1.32}$	0.318 $_{\pm 0.044}$	15.62 s	26.64 $_{\pm 1.52}$	0.240 $_{\pm 0.029}$	15.49 s	25.34 $_{\pm 1.52}$	0.284 $_{\pm 0.038}$	18.19 s	N/A	N/A	N/A	N/A	N/A	N/A
IDPG [14] [CVPR2024]	27.20 $_{\pm 1.68}$	0.326 $_{\pm 0.047}$	8.42 s	25.51 $_{\pm 1.22}$	0.411 $_{\pm 0.054}$	8.53 s	27.47 $_{\pm 1.83}$	0.313 $_{\pm 0.035}$	8.46 s	26.02 $_{\pm 1.56}$	0.354 $_{\pm 0.048}$	9.77 s	22.37 $_{\pm 1.29}$	0.261 $_{\pm 0.033}$	8.48 s	N/A	N/A	N/A
ProjDiff [50] [NearFPS2024]	27.09 $_{\pm 1.53}$	0.242 $_{\pm 0.031}$	8.13 s	25.73 $_{\pm 1.11}$	0.336 $_{\pm 0.046}$	8.15 s	27.91 $_{\pm 1.78}$	0.238 $_{\pm 0.029}$	8.11 s	N/A	N/A	N/A	22.06 $_{\pm 1.26}$	0.248 $_{\pm 0.033}$	8.32 s	N/A	N/A	N/A
DDNM [41] [ICLR2023]	27.45 $_{\pm 1.12}$	0.245 $_{\pm 0.035}$	8.03 s	25.52 $_{\pm 1.48}$	0.324 $_{\pm 0.047}$	8.06 s	27.69 $_{\pm 0.25}$	0.237 $_{\pm 0.027}$	7.97 s	N/A	N/A	N/A	21.66 $_{\pm 1.35}$	0.265 $_{\pm 0.041}$	8.20 s	+2.623%	-9.944%	-3.044%
+ EquS* [43] [WACV2026]	27.63 $_{\pm 0.25}$	0.242 $_{\pm 0.033}$	8.00 s	25.56 $_{\pm 1.37}$	0.327 $_{\pm 0.045}$	8.12 s	27.72 $_{\pm 1.87}$	0.234 $_{\pm 0.025}$	8.01 s	N/A	N/A	N/A	22.14 $_{\pm 1.25}$	0.246 $_{\pm 0.038}$	8.21 s			
+ LEADer [Ours]	27.71 $_{\pm 1.92}$	0.231 $_{\pm 0.029}$	7.74 s	25.63 $_{\pm 1.28}$	0.306 $_{\pm 0.038}$	7.86 s	27.81 $_{\pm 1.65}$	0.207 $_{\pm 0.022}$	7.66 s	N/A	N/A	N/A	23.54 $_{\pm 1.36}$	0.223 $_{\pm 0.030}$	8.02 s			
DDPG [14] [CVPR2024]	27.41 $_{\pm 1.86}$	0.255 $_{\pm 0.031}$	8.43 s	25.55 $_{\pm 1.42}$	0.354 $_{\pm 0.044}$	8.17 s	27.73 $_{\pm 1.74}$	0.205 $_{\pm 0.026}$	8.24 s	25.94 $_{\pm 1.58}$	0.249 $_{\pm 0.033}$	8.79 s	20.69 $_{\pm 1.27}$	0.238 $_{\pm 0.037}$	8.56 s	+1.079%	-4.914%	-3.169%
+ EquS [43] [WACV2026]	27.44 $_{\pm 1.73}$	0.252 $_{\pm 0.029}$	8.49 s	25.63 $_{\pm 1.35}$	0.351 $_{\pm 0.041}$	8.23 s	27.73 $_{\pm 1.61}$	0.203 $_{\pm 0.024}$	8.31 s	25.99 $_{\pm 1.46}$	0.247 $_{\pm 0.031}$	8.86 s	20.71 $_{\pm 1.18}$	0.255 $_{\pm 0.034}$	8.56 s			
+ LEADer [Ours]	27.62 $_{\pm 1.52}$	0.243 $_{\pm 0.030}$	8.18 s	25.75 $_{\pm 1.21}$	0.339 $_{\pm 0.036}$	7.95 s	28.02 $_{\pm 1.43}$	0.194 $_{\pm 0.021}$	8.00 s	26.19 $_{\pm 1.32}$	0.236 $_{\pm 0.028}$	8.51 s	21.07 $_{\pm 1.09}$	0.245 $_{\pm 0.030}$	8.21 s			
SITCOM [2] [ICML2025]	27.26 $_{\pm 1.32}$	0.186 $_{\pm 0.026}$	154.64 s	25.35 $_{\pm 1.21}$	0.232 $_{\pm 0.038}$	163.50 s	27.40 $_{\pm 0.45}$	0.236 $_{\pm 0.039}$	142.36 s	28.65 $_{\pm 0.34}$	0.189 $_{\pm 0.036}$	140.87 s	N/A	N/A	N/A	+0.809%	-4.828%	-5.470%
+ EquS [43] [WACV2026]	27.30 $_{\pm 1.30}$	0.181 $_{\pm 0.024}$	151.88 s	25.36 $_{\pm 1.19}$	0.228 $_{\pm 0.036}$	160.42 s	27.38 $_{\pm 0.44}$	0.232 $_{\pm 0.037}$	139.93 s	28.66 $_{\pm 0.33}$	0.185 $_{\pm 0.034}$	138.51 s	N/A	N/A	N/A			
+ LEADer [Ours]	27.56 $_{\pm 1.28}$	0.176 $_{\pm 0.022}$	145.97 s	25.53 $_{\pm 1.16}$	0.223 $_{\pm 0.033}$	154.36 s	27.59 $_{\pm 0.43}$	0.226 $_{\pm 0.035}$	134.68 s	28.86 $_{\pm 0.32}$	0.178 $_{\pm 0.031}$	133.42 s	N/A	N/A	N/A			
PIRP F ₀ G. [47] [ORS2026]	24.86 $_{\pm 1.34}$	0.151 $_{\pm 0.024}$	8.55 s	25.07 $_{\pm 1.32}$	0.290 $_{\pm 0.041}$	8.56 s	26.17 $_{\pm 1.28}$	0.150 $_{\pm 0.022}$	8.69 s	24.81 $_{\pm 1.67}$	0.209 $_{\pm 0.035}$	9.83 s	23.68 $_{\pm 1.45}$	0.217 $_{\pm 0.031}$	8.16 s	+1.487%	-2.726%	-4.010%
+ EquS [43] [WACV2026]	24.91 $_{\pm 1.21}$	0.153 $_{\pm 0.021}$	8.58 s	25.02 $_{\pm 1.39}$	0.294 $_{\pm 0.038}$	8.61 s	26.11 $_{\pm 1.16}$	0.155 $_{\pm 0.020}$	8.73 s	24.76 $_{\pm 1.48}$	0.211 $_{\pm 0.032}$	9.88 s	23.71 $_{\pm 1.33}$	0.220 $_{\pm 0.029}$	8.19 s			
+ LEADer [Ours]	25.37 $_{\pm 1.05}$	0.146 $_{\pm 0.018}$	8.29 s	25.34 $_{\pm 1.12}$	0.278 $_{\pm 0.030}$	8.32 s	26.28 $_{\pm 0.98}$	0.152 $_{\pm 0.019}$	8.41 s	25.24 $_{\pm 1.21}$	0.201 $_{\pm 0.028}$	9.63 s	24.19 $_{\pm 1.10}$	0.209 $_{\pm 0.026}$	7.43 s			

analytical upper bound:

$$\|x^{\text{Dense}} - x^{\text{LEADer}}\|_2 \leq C \cdot B \cdot \frac{e^{LT} - 1}{L}, \quad (23)$$

where $C > 0$ is a constant, T is the total samplings steps of the reverse process, and L is the Lipschitz constant.

In the discrete ODE solution process, due to Lipschitz continuity, the local truncation error accumulates exponentially over time steps, yielding the e^{LT} term. By introducing the uncertainty trace U_t to adjust the sampling step size, SATP constrains the local error at each step within the predefined tolerance B . This state-aware local error control ensures that the global accumulated error remains bounded under non-uniform sampling, thereby guaranteeing the convergence of the accelerated generation trajectory.

5 Experiments

5.1 Experimental Setup

We conduct experiments on five typical IR tasks: 4 \times Super-Resolution (SR) with a bicubic downsampler ($\sigma_e = 0$ and $\sigma_e = 0.05$), Gaussian deblurring ($\sigma_e = 0.05$), motion deblurring ($\sigma_e = 0.05$), and Compressed Sensing (CS) using a Walsh-Hadamard sampling matrix with a 0.25 compression ratio. For validation, we use two standard image restoration benchmarks, i.e., CelebA-HQ 1K [23] and ImageNet 1K [30], both at 256×256 resolution. For a fair comparison, all methods utilize the same pre-trained DDM denoisers, i.e., one trained on CelebA-HQ [25] and one trained on ImageNet [10]. Sampling steps are uniformly set to $T = 100$ for all methods, except for DPS [5] ($T = 1000$) and SITCOM [2]. Note that due to the step-skipping capability of SATP, our total number of Function

Evaluations (NFEs) is less than 100. All experiments were performed on an NVIDIA GeForce RTX 3090 GPU.

Evaluation Metrics. We employ three prevalent metrics to evaluate the proposed method on image restoration tasks,


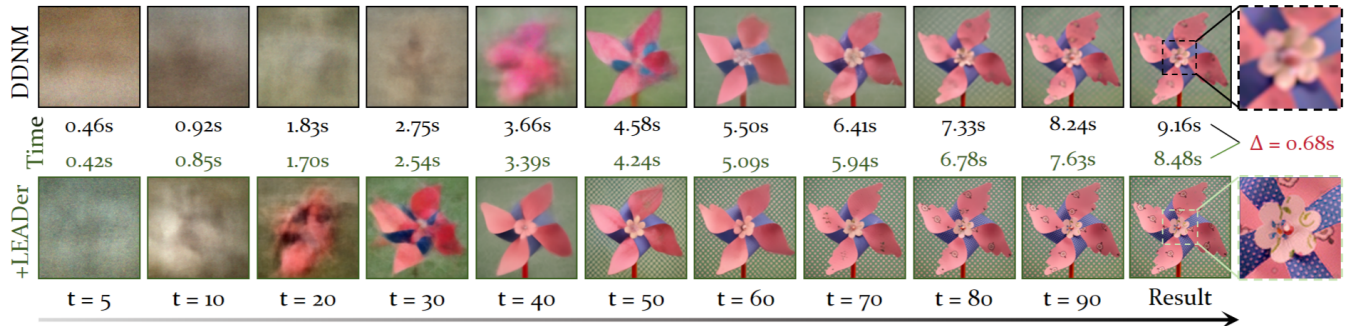


Figure 3: Qualitative comparisons on five image restoration tasks.

Figure 4: Comparison of $\hat{x}_{0|t}$ between DDNM and LEADer at different NFEs.

significantly reducing LPIPS by 7.13% to 12.78%. Notably, unlike EquS [43], LEADer leverages a state-aware trajectory pruning strategy to improve generation quality while achieving substantial sampling speedups, saving 3.04% to 8.42% of sampling time on average. These results demonstrate that our local epistemic uncertainty quantification directly strengthens the adaptive adjustment for complex local structures while significantly boosting inference efficiency.

Qualitative Results. To assess the visual restoration capability of LEADer under complex degradations, we compare the visual results of various baseline methods before and after its integration, as shown in Figure 3. When confronted with complex localized degradations, baseline methods often suffer from detail loss or over-smoothing due to their fixed prior weights. By integrating LEADer, the model can adaptively modulate the prior strength, which effectively suppresses artifacts while recovering sharper edges and more realistic textures, thus leading to a substantial improvement in overall visual fidelity.

To further illustrate the effect of the uncertainty-based active sampling strategy, we visualize the restoration trajectories of $\hat{x}_{0|t}$ across different steps on the $4\times$ super-resolution task, as shown in Figure 4. Due to its fixed step size, the baseline DDNM recovers details through a slow and uniform process, where both intermediate states and final outputs often exhibit high uncertainty. In contrast, LEADer perceives the current state and dynamically adjusts the sampling pace, which enables earlier convergence to clear local textures at lower steps. This enhanced sensitivity to detail not only improves data consistency but also allows the SATP strategy to safely skip steps, thereby reducing inference time.

5.3 Ablation Study

Impact of Core Strategies. To validate the effectiveness of the key strategies in the LEADer framework, we conduct ablation studies on three image restoration tasks using the DDNM [41] baseline. As shown in Table 2, introducing UCPM alone significantly improves

Table 2: Ablation study results of each component in LEADER.

DDNM			Bicub. SR $\times 4$ ($\sigma_e = 0.05$)			Gaussian Deb. ($\sigma_e = 0.05$)			CS 25% ($\sigma_e = 0.05$)		
UCPM	SATP	T	PSNR \uparrow	LPIPS \downarrow	Time \downarrow	PSNR \uparrow	LPIPS \downarrow	Time \downarrow	PSNR \uparrow	LPIPS \downarrow	Time \downarrow
×	×	100	29.19	0.082	3.87 s	30.60	0.072	3.90 s	26.14	0.124	3.83 s
✓	×	100	29.69	0.075	3.96 s	30.90	0.061	4.01 s	27.14	0.090	3.97 s
×	×	93	28.85	0.098	3.67 s	30.31	0.089	3.83 s	25.79	0.134	3.75 s
×	✓	100	29.17	0.084	3.66 s	30.58	0.077	3.80 s	26.10	0.125	3.73 s
✓	×	93	29.35	0.090	3.65 s	30.58	0.074	3.75 s	26.74	0.106	3.70 s
✓	✓	100	29.69	0.077	3.63 s	30.86	0.063	3.75 s	27.06	0.094	3.69 s


Table 3: Hyperparameter sensitivity analysis on the information loss budget B .

ImageNet 1K Setting	Bicub. SR $\times 4$ ($\sigma_e = 0$)			Motion Deb. ($\sigma_e = 0.05$)		
	PSNR \uparrow	LPIPS \downarrow	Time \downarrow	PSNR \uparrow	LPIPS \downarrow	Time \downarrow
$B = 0$	25.41	0.143	8.56	25.31	0.196	9.83
$B = 0.001$	25.41	0.143	8.56	25.30	0.196	9.83
$B = 0.005$	25.38	0.144	8.45	25.26	0.199	9.74
$B = 0.01$	25.37	0.146	8.29	25.24	0.201	9.63
$B = 0.05$	24.83	0.161	4.54	24.71	0.226	5.21
$B = 0.1$	23.96	0.187	2.76	23.85	0.258	3.08

the restoration quality, albeit with a slight increase in sampling time. Conversely, employing SATP alone effectively reduces inference time with virtually no quality degradation. When combined, they exhibit excellent complementarity, achieving superior restoration quality while simultaneously reducing the overall sampling time. This indicates that the effective local guidance provided by UCPM enhances the stability of reverse sampling, thereby enabling SATP to perform step-skipping more safely. Furthermore, to demonstrate the advantage of the SATP strategy over directly reducing sampling steps, we conduct a step ablation study. The results show that, under the condition of roughly equivalent sampling time, our active sampling strategy demonstrates significant robustness. These experimental results demonstrate that both UCPM and SATP are indispensable, yielding clear synergy and mutual enhancement across spatiotemporal dimensions.

Impact of Sampling Steps. To evaluate the robustness of our method across varying sampling steps, we compare the performance of DDNM [41] and DDPG [14] with and without LEADER at $T \in \{10, 20, 50, 100\}$. As shown in Figure 5, introducing LEADER consistently improves image quality across all settings, thus demonstrating the effectiveness of our UCPM. In terms of efficiency, SATP consistently reduces inference time. However, under extremely low step counts, the runtime reduction becomes marginal. This is because the overall epistemic uncertainty remains excessively high under such sparse sampling; consequently, SATP adaptively performs fewer skip operations to preserve restoration quality. Overall, these results verify that LEADER exhibits strong robustness across different sampling steps.

Impact of B . The information loss budget B in Equation (18) governs the trade-off between restoration quality and computational efficiency. A larger B tolerates higher local truncation errors, thereby encouraging SATP to execute more aggressive trajectory pruning. To investigate the impact of B , we conduct a sensitivity analysis using PIRP $F_{IB}G_-$ [47] as the baseline on the ImageNet 1K dataset for the Bicubic SR $\times 4$ ($\sigma_e = 0$) and Motion Deblurring ($\sigma_e = 0.05$) tasks. As shown in Table 3, a conservative budget ($B = 0.001$) yields quality improvements while restricting acceleration. Conversely, an

**Figure 5: Ablation study on different sampling steps. (Dataset: CelebA-HQ 1K, Task: Gaussian Deb. & $\sigma = 0.05$)****Table 4: Memory consumption comparison across five tasks.**

Methods	Bicub. SR $\times 4$ ($\sigma_e = 0$)	Bicub. SR $\times 4$ ($\sigma_e = 0.05$)	Gaussian Deb. ($\sigma_e = 0.05$)	Motion Deb. ($\sigma_e = 0.05$)	CS 25% ($\sigma_e = 0.05$)
DDNM [41]	2507 MB	2513 MB	2509 MB	N/A	2541 MB
+ LEADER	2509 MB	2516 MB	2515 MB	N/A	2543 MB
DDPG [14]	2521 MB	2524 MB	2523 MB	2959 MB	2564 MB
+ LEADER	2526 MB	2526 MB	2530 MB	2965 MB	2567 MB
SITCOM [2]	15831 MB	15863 MB	15912 MB	18478 MB	N/A
+ LEADER	15842 MB	15872 MB	15923 MB	18485 MB	N/A
PIRP $F_{IB}G_-$ [47]	2513 MB	2520 MB	2515 MB	2948 MB	2557 MB
+ LEADER	2516 MB	2521 MB	2521 MB	2952 MB	2561 MB

overly aggressive budget ($B = 0.1$), despite significantly reducing inference time, leads to notable performance degradation. Based on these observations, we set $B = 0.01$ as the default configuration for our SATP strategy to optimally balance restoration fidelity and sampling efficiency.

5.4 Memory Consumption

To evaluate the practical overhead of LEADER as a plug-and-play module, we compare the GPU memory consumption of various baseline methods before and after its integration, as shown in Table 4. Evaluations across five tasks on the ImageNet 1K dataset reveal that despite introducing additional computations, LEADER incurs a marginal overall GPU memory increase of merely 0.07% to 0.28%. This indicates that LEADER can robustly enhance both the restoration quality and sampling efficiency of existing methods at an almost negligible memory cost.

6 Conclusion

In this paper, we identify a fundamental limitation of existing zero-shot DMIR methods, i.e., their reliance on fixed prior constraints and uniform sampling schedules, which often leads to local structural distortions and redundant computations. To address this, we propose LEADER, an active diffusion sampling framework guided by local epistemic uncertainty. Spatially, we leverage pixel-level uncertainty to dynamically modulate prior strength, which effectively balances detail preservation and artifact suppression. Temporally, we quantify the stability of the generation process via uncertainty traces and adaptively prune trajectories within deterministic error bounds. Both theoretical analysis and extensive experiments demonstrate that LEADER can be seamlessly integrated as a plug-and-play module to consistently improve restoration quality and sampling efficiency, while preserving strict data consistency. In the future, it would be also interesting to investigate the application of this uncertainty-based active sampling mechanism to diffusion-model-based generative tasks.

References

- [1] Yasar Utku Alçalar and Mehmet Akçakaya. 2024. Zero-Shot Adaptation for Approximate Posterior Sampling of Diffusion Models in Inverse Problems. In *European Conference on Computer Vision (Lecture Notes in Computer Science, Vol. 15141)*. Springer, 444–460.
- [2] Ismail Alkhouri, Shijun Liang, Cheng-Han Huang, Jimmy Dai, Qing Qu, Saiprasad Ravishankar, and Rongrong Wang. 2025. SITCOM: Step-wise Triple-Consistent Diffusion Sampling For Inverse Problems. In *International Conference on Machine Learning*, Vol. 267. PMLR / OpenReview.net.
- [3] Yochai Blau and Tomer Michaeli. 2018. The Perception-Distortion Tradeoff. In *IEEE Conference on Computer Vision and Pattern Recognition*. Computer Vision Foundation / IEEE Computer Society, 6228–6237.
- [4] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. 2022. Simple Baselines for Image Restoration. In *European Conference on Computer Vision (Lecture Notes in Computer Science, Vol. 13667)*. Springer, 17–33.
- [5] Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul Ye. 2023. Diffusion Posterior Sampling for General Noisy Inverse Problems. In *International Conference on Learning Representations*. OpenReview.net.
- [6] Hyungjin Chung, Jeongsol Kim, and Jong Chul Ye. 2023. Direct Diffusion Bridge using Data Consistency for Inverse Problems. In *Annual Conference on Neural Information Processing Systems*.
- [7] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. 2022. Improving Diffusion Models for Inverse Problems using Manifold Constraints. In *Annual Conference on Neural Information Processing Systems*.
- [8] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. 2023. Diffusion Models in Vision: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 9 (2023), 10850–10869.
- [9] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen O. Egiazarian. 2007. Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering. *IEEE Transactions on Image Processing* 16, 8 (2007), 2080–2095.
- [10] Prafulla Dhariwal and Alexander Quinn Nichol. 2021. Diffusion Models Beat GANs on Image Synthesis. In *Annual Conference on Neural Information Processing Systems*. 8780–8794.
- [11] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. 2016. Image Super-Resolution Using Deep Convolutional Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 2 (2016), 295–307.
- [12] Bradley Efron. 2011. Tweedie’s formula and selection bias. *J. Amer. Statist. Assoc.* 106, 496 (2011), 1602–1614.
- [13] Ben Fei, Zhaoyang Lyu, Liang Pan, Junzhe Zhang, Weidong Yang, Tianyue Luo, Bo Zhang, and Bo Dai. 2023. Generative Diffusion Prior for Unified Image Restoration and Enhancement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 9935–9946.
- [14] Tomer Garber and Tom Tiram. 2024. Image Restoration by Denoising Diffusion Models with Iteratively Preconditioned Guidance. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 25245–25254.
- [15] Shuhang Gu, Lei Zhang, Wangmeng Zuo, and Xiangchu Feng. 2014. Weighted Nuclear Norm Minimization with Application to Image Denoising. In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society, 2862–2869.
- [16] Linchao He, Hongyu Yan, Mengting Luo, Hongjie Wu, Kunming Luo, Wang Wang, Wenchao Du, Hu Chen, Hongyu Yang, Yi Zhang, et al. 2023. Fast and stable diffusion inverse solver with history gradient update. *CoRR abs/2307.12070* (2023).
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising Diffusion Probabilistic Models. In *Annual Conference on Neural Information Processing Systems*.
- [18] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. 2022. Elucidating the Design Space of Diffusion-Based Generative Models. In *Annual Conference on Neural Information Processing Systems*.
- [19] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. 2022. Denoising Diffusion Restoration Models. In *Annual Conference on Neural Information Processing Systems*.
- [20] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P. Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. 2017. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society, 105–114.
- [21] Xiang Li, Soo Min Kwon, Ismail R. Alkhouri, Saiprasad Ravishankar, and Qing Qu. 2024. Decoupled Data Consistency with Diffusion Purification for Image Restoration. *CoRR abs/2403.06054* (2024).
- [22] Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. 2021. SwinIR: Image Restoration Using Swin Transformer. In *IEEE/CVF International Conference on Computer Vision Workshops*. IEEE, 1833–1844.
- [23] Ziwei Liu, Ping Luo, Xiaoogang Wang, and Xiaoou Tang. 2015. Deep Learning Face Attributes in the Wild. In *International Conference on Computer Vision*. IEEE Computer Society, 3730–3738.
- [24] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. 2022. DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps. In *Annual Conference on Neural Information Processing Systems*.
- [25] Andreas Lugmayr, Martin Danelljan, Andrés Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. 2022. RePaint: Inpainting using Denoising Diffusion Probabilistic Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 11451–11461.
- [26] Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. 2024. A Variational Perspective on Solving Inverse Problems with Diffusion Models. In *International Conference on Learning Representations*. OpenReview.net.
- [27] Michael T. McCann, Kyong Hwan Jin, and Michael Unser. 2017. Convolutional Neural Networks for Inverse Problems in Imaging: A Review. *IEEE Signal Processing Magazine* 34, 6 (2017), 85–95.
- [28] Mang Ning, Enver Sangineto, Angelo Porrello, Simone Calderara, and Rita Cucchiara. 2023. Input Perturbation Reduces Exposure Bias in Diffusion Models. In *International Conference on Machine Learning (Proceedings of Machine Learning Research)*. PMLR, 26245–26265.
- [29] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 10674–10685.
- [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *Int. J. Comput. Vis.* 115 (2015), 211–252.
- [31] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *International Conference on Machine Learning (JMLR Workshop and Conference Proceedings, Vol. 37)*, Francis R. Bach and David M. Blei (Eds.). JMLR.org, 2256–2265.
- [32] Jiaming Song, Chenlin Meng, and Stefano Ermon. 2021. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*. OpenReview.net.
- [33] Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. 2023. Pseudoinverse-Guided Diffusion Models for Inverse Problems. In *International Conference on Learning Representations*. OpenReview.net.
- [34] Yang Song and Stefano Ermon. 2019. Generative Modeling by Estimating Gradients of the Data Distribution. In *Annual Conference on Neural Information Processing Systems*. 11895–11907.
- [35] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. 2022. Solving Inverse Problems in Medical Imaging with Score-Based Generative Models. In *International Conference on Learning Representations*. OpenReview.net.
- [36] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*. OpenReview.net.
- [37] Charles M Stein. 1981. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics* (1981), 1135–1151.
- [38] Darshan Thaker, Abhishek Goyal, and René Vidal. 2024. Frequency-Guided Posterior Sampling for Diffusion-Based Image Restoration. *CoRR abs/2411.15295* (2024).
- [39] Chong Wang, Lanqing Guo, Zixuan Fu, Siyuan Yang, Hao Cheng, Alex C. Kot, and Bihan Wen. 2025. Reconciling Stochastic and Deterministic Strategies for Zero-shot Image Restoration using Diffusion Model in Dual. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. Computer Vision Foundation / IEEE, 23207–23216.
- [40] Ge Wang, Jong Chul Ye, Klaus Mueller, and Jeffrey A. Fessler. 2018. Image Reconstruction is a New Frontier of Machine Learning. *IEEE Transactions on Medical Imaging* 37, 6 (2018), 1289–1296.
- [41] Yinhuai Wang, Jiwen Yu, and Jian Zhang. 2023. Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model. In *International Conference on Learning Representations*. OpenReview.net.
- [42] Zhihao Wang, Jian Chen, and Steven C. H. Hoi. 2021. Deep Learning for Image Super-Resolution: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43, 10 (2021), 3365–3387.
- [43] Chenxu Wu, Qingpeng Kong, Peiang Zhao, Wendi Yang, Wenxin Ma, Fenghe Tang, Zihang Jiang, and S. Kevin Zhou. 2025. Equivariant Sampling for Improving Diffusion Model-based Image Restoration. *CoRR abs/2511.09965* (2025).
- [44] Hongjie Wu, Linchao He, Mingqin Zhang, Dongdong Chen, Kunming Luo, Mengting Luo, Jizhe Zhou, Hu Chen, and Jiancheng Lv. 2024. Diffusion Posterior Proximal Sampling for Image Restoration. In *ACM International Conference on Multimedia*. ACM, 214–223.
- [45] Hongjie Wu, Mingqin Zhang, Linchao He, Ji-Zhe Zhou, and Jiancheng Lv. 2025. Enhancing Diffusion Model Stability for Image Restoration via Gradient Management. In *ACM International Conference on Multimedia*. ACM, 10768–10777.
- [46] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. 2024. Diffusion Models: A Comprehensive Survey of Methods and Applications. *Comput. Surveys* 56, 4

1045	(2024), 105:1–105:39.		
1046	[47] Yang Yang, Xi Zhang, Jiaqi Zhang, and Lanling Zeng. 2026. Parameterized image restoration with diffusion and gradient priors. <i>Knowledge Based Systems</i> 338 (2026), 115488.		
1047			
1048	[48] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. 2022. Restormer: Efficient Transformer for High-Resolution Image Restoration. In <i>IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> . IEEE, 5718–5729.		
1049			
1050			
1051	[49] Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang Song. 2025. Improving Diffusion Inverse Problem Solving with Decoupled Noise Annealing. In <i>IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> . IEEE, 20895–20905.		
1052			
1053	[50] Jiawei Zhang, Jiaxin Zhuang, Cheng Jin, Gen Li, and Yuantao Gu. 2024. Unleashing the Denoising Capability of Diffusion Prior for Solving Inverse Problems. In <i>Annual Conference on Neural Information Processing Systems</i> .		
1054			
1055			
1056			
1057			
1058			
1059			
1060			
1061			
1062			
1063			
1064			
1065			
1066			
1067			
1068			
1069			
1070			
1071			
1072			
1073			
1074			
1075			
1076			
1077			
1078			
1079			
1080			
1081			
1082			
1083			
1084			
1085			
1086			
1087			
1088			
1089			
1090			
1091			
1092			
1093			
1094			
1095			
1096			
1097			
1098			
1099			
1100			
1101			
1102			
		[51] Kai Zhang, Yawei Li, Wangmeng Zuo, Lei Zhang, Luc Van Gool, and Radu Timofte. 2022. Plug-and-Play Image Restoration With Deep Denoiser Prior. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> 44, 10 (2022), 6360–6376.	1103
			1104
			1105
		[52] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. 2017. Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising. <i>IEEE Transactions on Image Processing</i> 26, 7 (2017), 3142–3155.	1106
			1107
		[53] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In <i>IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> . 586–595.	1108
			1109
		[54] Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc Van Gool. 2023. Denoising Diffusion Models for Plug-and-Play Image Restoration. In <i>IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> . IEEE, 1219–1229.	1110
			1111
			1112
			1113
			1114
			1115
			1116
			1117
			1118
			1119
			1120
			1121
			1122
			1123
			1124
			1125
			1126
			1127
			1128
			1129
			1130
			1131
			1132
			1133
			1134
			1135
			1136
			1137
			1138
			1139
			1140
			1141
			1142
			1143
			1144
			1145
			1146
			1147
			1148
			1149
			1150
			1151
			1152
			1153
			1154
			1155
			1156
			1157
			1158
			1159
			1160