# PLUG-IN IMAGE QUALITY CONTROL FOR POSTERIOR DIFFUSION SUPER-RESOLUTION

**Anonymous authors**Paper under double-blind review

000

001

002003004

010 011

012

013

014

016

018

019

021

024

025

026

027

028

029

031

033

037

038

040

041

042

043

044

046

047

051

052

#### **ABSTRACT**

Diffusion-based super-resolution (SR) has shown remarkable progress, mainly through prior-guided approaches that require explicit degradation models or semantic priors. While posterior diffusion SR avoids these assumptions by directly learning from LR-HR pairs, it still suffers from numerical errors during sampling and lacks plug-and-play mechanisms for quality control. We introduce the first plug-and-play framework for posterior diffusion SR, enabling pretrained models to support controllable quality without retraining. Our numerical analysis reveals that discretization errors are a key bottleneck in posterior SR. We prove that these errors can be equivalently expressed as gradients of KL divergence, unifying numerical error correction with image based classifier guidance. This provides a principled explanation of fidelity degradation and a new lens for posterior diffusion trajectories. In principle, these errors can be corrected to improve fidelity when reference supervision is available, offering a new theoretical understanding of posterior diffusion trajectories. In real-world SR, we further show that our image based guidance offers a controllable trade-off between fidelity and perception, delivering perceptual sharpness competitive with state-of-the-art prior-based models. Experiments confirm that our method consistently improves perceptual quality, while also validating the theoretical link between numerical errors and fidelity in posterior SR. These results position our work as a new direction for posterior diffusion models, bridging probabilistic analysis with practical deployment.

#### 1 Introduction

Image super-resolution (SR) is a fundamental problem in low-level vision, aiming to recover high-resolution (HR) images from their low-resolution (LR) counterparts. Classical approaches include regression-based methods Mehta et al. (2023); Wang et al. (2023a); Lim et al. (2017); Liang et al. (2021a) and adversarial training with GANs Ledig et al. (2017); Wang et al. (2021; 2018), but these methods still struggle with the intrinsic ill-posedness of SR.

Diffusion models Ho et al. (2020); Song et al. (2021b) have recently emerged as a powerful generative framework, leveraging a physics-driven probabilistic process that excels at reconstructing high-frequency details in ill-posed image restoration problems Saharia et al. (2022b); Wang et al. (2023c); Li et al. (2022a); Kawar et al. (2022). While diffusion models were initially criticized for their high computational cost, subsequent advances such as improved samplers Song et al. (2021a); Yue et al. (2023) and distillation-based acceleration Wu et al. (2024a); Dong et al. (2024); Chen et al. (2025; 2024c); Yue et al. (2025); Lin et al. (2024); Wu et al. (2024b); Yang et al. (2024) have reduced the sampling steps to single or few iterations, making runtime comparable to regression-based SR. Moreover, semantic priors such as CLIP have been employed to guide perceptual quality.

Despite these advances, most existing works adopt a *diffusion prior* paradigm: pretrained generative priors with explicit or learned degradation models Saharia et al. (2021); Li et al. (2022c). However, diffusion priors require an accurate likelihood formulation of the degradation process, which is often intractable beyond simplistic assumptions such as Gaussian blur or bicubic downsampling. While higher-order solvers Lu et al. (2022; 2023); Zheng et al. (2023) mitigate discretization error in the generative trajectory, they remain dependent on assumed likelihoods that do not align well with real-world degradations, especially those introduced by optical lenses.

n	=	Л
U	J	4
_	_	_

060

061

062

063

064

065 066

067

068

071

073 074

075

076

077

078

079

081

082

084

085

090

092

095

096

098

102

104

105

106

107

Table 1: Prior vs. Posterior Diffusion SR

Aspect	Diffusion Prior SR	Posterior Diffusion SR			
Training	Learn prior $p(x)$ + likelihood $p(y x)$	Directly learn $p(x y)$ from LR–HR pairs			
Assumption	need explicit degradation model. Bicubic, Gaussian)	Covered by Dataset			
Plug-and-play	CLIP Radford et al. (2021)	Ours: first plug-in model			
Distillation	Actively used	Not Active			

In contrast, *posterior learning* approaches define the degradation process directly through LR-HR training pairs Cai et al. (2019a); Wang et al. (2021); Bhat et al. (2021). This formulation bypasses the need for explicit likelihood modeling and naturally inherits the legacy of regression-based SR datasets. However, posterior diffusion models remain underexplored: unlike diffusion priors, no plug-and-play modules for image quality control have been established on pretrained posterior models.

From a numerical analysis perspective, diffusion samplers typically employ first-order forward discretization, which is computationally efficient but induces significant numerical error. Existing efforts Lu et al. (2022; 2023); Zheng et al. (2023) target discretization in diffusion priors, but posterior models have not been systematically analyzed in this regard. Notably, Yue et al. (2023) adopts a single-order solver, which remains susceptible to discretization error. For general image generation, Li et al. Li & van der Schaar (2024) analyzed cumulative error between forward and backward processes, but the role of discretization error in SR has yet to be explicitly studied.

Our contribution. The contributions of this work are summarized as follows:

- We propose the first plug-and-play module for posterior diffusion SR, enabling pretrained models to support image quality control without retraining.
- We theoretically derive that discretization errors in first-order posterior models can be corrected by incorporating second-order differential terms, providing a principled explanation of fidelity degradation.
- We further propose an *image-based classifier guidance formulation* for posterior diffusion and prove that it is mathematically equivalent to the second-order correction term, thereby unifying numerical analysis with guidance-based conditioning.
- While explicit fidelity correction requires HR references and is thus of limited use in blind SR, we show that the same formulation can be adapted in a sign-flipped manner to enhance perceptual quality in reference-free scenarios.
- This dual view—fidelity-oriented correction under supervision and perception-oriented control without supervision—establishes a unified framework for posterior SR.
- Extensive analysis and experiments validate that our approach bridges the gap between numerical error theory and real-world perceptual enhancement in posterior learning models.

### 2 Related Works

**Image Super-Resolution** Since the emergence of deep neural network-based image super-resolution Liang et al. (2021a); Zhang et al. (2021); Dong et al. (2011); Gu et al. (2015); Zamir et al. (2022), significant advancements have been made in this field. Generative adversarial networks (GANs) Goodfellow et al. (2014) have played a crucial role in enhancing details in super-resolution outputs Wang et al. (2021; 2018). By leveraging adversarial learning, GANs improve texture sharpness and realism in super-resolved images. Recently, diffusion models Saharia et al. (2022b); Li et al. (2022a); Kawar et al. (2022); Yue et al. (2023); Wang et al. (2023c) have emerged as a powerful alternative for image super-resolution. Originally developed for novel image synthesis, these models have demonstrated exceptional capabilities in texture restoration and fine-detail reconstruction. Unlike GANs, diffusion models progressively refine images from noise, providing a more stable and controlled SR framework. Studies have shown that diffusion-based SR outperforms GAN-based approaches, particularly in preserving high-frequency details and reducing artifacts Rombach et al. (2022a). Furthermore, developments in VQGAN Esser et al. (2021) and stable diffusion models Rombach et al. (2022a) have significantly influenced image synthesis Saharia et al. (2022a), expanding beyond SR into text-conditioned image generation. These advancements have played a pivotal role in integrating SR capabilities within broader generative frameworks.

Performance Enhancement Diffusion models for image super-resolution face an inherent challenge of high computational overhead Saharia et al. (2022b); Ho et al. (2020). Various efforts have been made to address this issue Song et al. (2021a); Lu et al. (2022; 2023); Zheng et al. (2023); Wu et al. (2024a); Wang et al. (2024a). Starting with DDIM, subsequent works on DPM solvers have significantly reduced the number of iteration steps from thousands to just tens. Residual-based approaches in image synthesis Liu et al. (2024) and image super-resolution Yue et al. (2023) have demonstrated stable performance even with fewer than ten steps. More recently, techniques such as model distillation Wang et al. (2024a) and single-step SR Wu et al. (2024a) have emerged, further pushing the boundaries of efficiency in diffusion-based SR.

Error Estimation Reducing the number of iteration steps inherently increases step intervals in finite difference methods, making diffusion models more susceptible to numerical errors Strang (2007); Peter & Eckhard (1992); Hochbruck & Ostermann (2005). Consequently, research has been conducted to analyze and mitigate numerical errors and their effects Li & van der Schaar (2024); Li et al. (2024). Trajectory analysis Chen et al. (2024b;a) is also effective for detailed analysis of numerical errors. In Li & van der Schaar (2024), the authors investigated errors in pretrained diffusion models by analyzing discrepancies between forward and reverse processes. They defined \*\*modular errors\*\* between these processes and extended their analysis to \*\*cumulative errors\*\* throughout the entire process. Similarly, Li et al. (2024) examined how iterative inference steps exacerbate exposure bias due to training-inference discrepancies and proposed a method to mitigate this issue without requiring DPM retraining.

High Order Differential Equation Solver Diffusion based super resolution initially began by learning the posterior distribution.Li et al. (2022a); Saharia et al. (2022c). However, the diffusion posterior sampling(DPS) approach has gained attention for its efficiency, as it enables sampling using a pre-trained prior without requiring additional training costs. Most existing higher-order ODE/SDE solvers (Heun, RK2/3, DPM-Solver-2/3, etc.) have been developed in the context of DDPM-like models, where the prior is learned and the sampling process is derived from a score-based or probability-flow ODE interpretation Song et al. (2021a); Karras et al. (2022); Lu et al. (2022; 2023); Zheng et al. (2023); Liu et al. (2022). These works typically apply numerical integration of a pre-trained prior distribution, thus enabling explicit higher-order schemes for fast sampling. However, compared to the DPS approach, which reuses a pre-trained model trained with a distribution based loss, the posterior learning using pixel-based loss is still considered advatangeous in terms of final image quality. Yue et al. (2023) Thus, our setting focuses on a **learned-posterior** approach rather than a learned-prior approach, where the solver is based on the first-order Euler-type update. The higher-order solvers introduced in prior-based diffusion works may not be readily applicable to the learned posterior method without re-architecting the model to fit a score-based ODE sampling paradigm.

## 3 PLUG AND PLAY FOR POSTERIOR DIFFUSION MODEL

## 3.1 BACKGROUND FOR NUMERICAL ACCURACY

A Taylor series expansion can be written as

$$x(t+h) = x(t) + h \, \nabla x(t) \; + \; \frac{1}{2} \, h^2 \, \nabla^2 x(t) \; + \; \mathcal{O}(h^3).$$

In a first-order difference scheme, we usually adopt

$$x(t+h) = x(t) + h \nabla x(t) + DE,$$

where DE is the leading-order discretization error, whose dominant term is  $\frac{1}{2}\,h^2\,\nabla^2x(t)$ . If we explicitly include this second-order term in the difference scheme, the accuracy increases. Most existing diffusion models, however, only consider a first-order discretization, whereas in this paper, we adopt a second-order discretization that accounts for  $\frac{1}{2}\,h^2\,\nabla^2x(t)$ . Higher-order discretization methods and their errors have also been analyzed in Lu et al. (2022; 2023); Zheng et al. (2023). While conventional formulas rely on starting the restoration from white Gaussian noise (thus requiring many diffusion steps), recent works on one-step or few-step diffusion models employ large step intervals, which can cause large discretization errors. Therefore, it becomes necessary to correct such discretization errors.

**DPM-Solver Revisited.** Consider a random variable  $x_0$  sampled from  $q_0(x_0)$ . The forward SDE on the interval [0, T] is

$$d\mathbf{x}_t = f(t) \mathbf{x}_t dt + g(t) d\mathbf{w}_t, \quad \mathbf{x}_0 \sim q_0(\mathbf{x}_0),$$

where  $\mathbf{w}_t \in \mathbb{R}^D$  is a standard Wiener process. Song et al. (2021b) show that its reverse process from T to 0, given the marginal  $q_T(x_T)$ , is

$$d\mathbf{x}_t = [f(t)\mathbf{x}_t - g^2(t)\nabla_{\mathbf{x}}\log p_t(\mathbf{x}_t)]dt + g(t)d\bar{\mathbf{w}}_t, \quad \mathbf{x}_T \sim p_T(\mathbf{x}_T),$$

where  $\bar{\mathbf{w}}_t$  is a standard Wiener process in reverse time. From the above, Lu et al. (2022; 2023) transform it into a probability flow ODE for faster sampling:

$$d\mathbf{x}_t = \left[ f(t) \, \mathbf{x}_t \, - \, \frac{1}{2} g^2(t) \, \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t) \right] dt. \tag{1}$$

To estimate the score function  $\nabla_x \log q_t(x_t)$ , DPMs use a neural network  $\epsilon_\theta(x_t, t)$ . The parameter  $\theta$  is optimized by minimizing

$$\mathcal{L}(\theta) = \int_0^T \mathbb{E}_{q_t(\mathbf{x}_t)} \| \boldsymbol{\epsilon}_{\theta}(\mathbf{x}_t, t) + \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t) \|_2^2 dt,$$

yielding

$$d\mathbf{x}_t = f(t)\mathbf{x}_t + g^2(t)\boldsymbol{\epsilon}_{\theta}(\mathbf{x}_t, t) dt.$$

We can generate samples by numerically solving this ODE from T down to 0. Given an initial value  $x_s$  at time s>0, Lu et al. (2022) show that the solution  $x_t$  for  $t\in[0,s]$  can be written in integral form as

$$\mathbf{x}_t = \frac{\alpha_t}{\alpha_s} \mathbf{x}_s + \alpha_t \int_{\lambda_s}^{\lambda_t} e^{-\lambda} \, \epsilon_{\theta}(\mathbf{x}_{\lambda}) \, d\lambda, \tag{2}$$

which in discrete form approximates

$$\mathbf{x}_{t} = \frac{\alpha_{t}}{\alpha_{s}} \mathbf{x}_{s} - \alpha_{t} \sum_{n=0}^{k-1} \epsilon_{\theta}^{(n)}(\mathbf{x}_{\lambda_{t}}) \int_{\lambda_{s}}^{\lambda_{t}} F(\lambda) d\lambda, \tag{3}$$

where  $F(\lambda)$  collects certain integral factors (see Lu et al. (2022)), and  $\epsilon_{\theta}^{(n)}(\cdot)$  is the n-th order derivative for  $n \leq k-1$ . For  $k \geq 2$ , they approximate these derivatives using intermediate points (Runge–Kutta). Although evaluating the integral more finely can reduce the discretization error, in practice it is typically done via finite-difference approaches like Runge–Kutta Hochbruck & Ostermann (2005). In our paper, however, we propose directly computing the second-order derivatives by taking gradients of the neural network model.

#### 3.2 Background for Inverse Problem Solver

Super-resolution (SR) is an ill-posed inverse problem, where the goal is to recover the high-resolution (HR) image  $\mathbf{x}$  from its degraded low-resolution (LR) observation  $\mathbf{y}$ . The forward measurement process can be expressed as  $\mathbf{y} = A(\mathbf{x}) + \mathbf{n}$ ,  $\mathbf{y}, \mathbf{n} \in \mathbb{R}^n$ ,  $\mathbf{x} \in \mathbb{R}^d$  where  $A(\cdot) : \mathbb{R}^d \to \mathbb{R}^n$  is the degradation operator (e.g., bicubic downsampling, blur, or camera pipeline), and  $\mathbf{n}$  denotes measurement noise. From a Bayesian perspective, the posterior distribution is given by

$$p(\mathbf{x}|\mathbf{y}) = p(\mathbf{y}|\mathbf{x}) p(\mathbf{x})/p(\mathbf{y})$$

where  $p(\mathbf{x})$  the prior over natural images,  $p(\mathbf{y}|\mathbf{x})$  the likelihood for degradation,  $p(\mathbf{x}|\mathbf{y})$  the posterior. In **Diffusion Prior** methods Saharia et al. (2021); Li et al. (2022b); Wang et al. (2023c); Wu et al. (2024b), the diffusion model serves as the image prior  $p(\mathbf{x})$ . The likelihood  $p(\mathbf{y}|\mathbf{x})$  is typically assumed to be a Gaussian degradation model:  $p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}; A(\mathbf{x}), \sigma^2 I)$ . Conditioning the diffusion process on  $\mathbf{y}$  yields a posterior-guided reverse ODE:

$$d\mathbf{x}_t = \left[ f(t) \,\mathbf{x}_t \, - \, \frac{1}{2} g^2(t) \, \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t | \mathbf{y}) \right] dt, \tag{4}$$

where the score term  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t|\mathbf{y})$  incorporates the assumed likelihood. Thus, the accuracy of prior-based SR crucially depends on the correctness of the degradation model.

**Posterior learning** approaches Yue et al. (2023) directly learn the conditional distribution  $p(\mathbf{x}|\mathbf{y})$  from paired LR–HR data, bypassing the need for an explicit likelihood model. In this case, the degradation process  $\mathbf{A}(\cdot)$  does not need to be analytically specified; it is implicitly encoded through the training dataset. The training objective is to maximize the conditional log-likelihood:  $\mathcal{L}(\theta) = \mathbb{E}_{(\mathbf{x},\mathbf{y})\sim\mathcal{D}}\big[\log p_{\theta}(\mathbf{x}|\mathbf{y})\big]$ . This makes posterior approaches attractive for real-world SR, where the degradation is complex or unknown. However, posterior SR is sensitive to discretization errors during sampling and currently lacks plug-and-play quality control mechanisms, which are more natural in prior-based frameworks.

#### 3.3 Theoretical Foundation

Error estimation in diffusion models has been addressed by Lu et al. (2022; 2023); Zheng et al. (2023); Li & van der Schaar (2024); Li et al. (2024). Most prior work focuses on measuring the difference between forward and reverse processes or errors from sampling intervals but their restoration has been less studied. Our interest lies in the discretization errors of finite-difference methods Strang (2007) and SDEs Peter & Eckhard (1992). Below, we define how we measure discretization error and Kullback–Leibler(KL)-based divergence error in our diffusion model for image restoration.

Assumption 1. (Definition of Error in the Diffusion Path) In Li & van der Schaar (2024), they defined the error in path of diffusion trajectory. The modular error  $\mathbf{E}_t^{\mathrm{modular}}$  measures the accuracy that every module maps its input to the output.  $\mathbf{E}_t^{\mathrm{modular}} = \mathbb{E} \big[ D_{\mathrm{KL}} \big( p_{\theta}(\mathbf{x}_{t-1} \mid \mathbf{x}_t) \, \| \, q(\mathbf{x}_{t-1} \mid \mathbf{x}_t) \big) \big]$  The cumulative error  $\mathbf{E}_t^{\mathrm{cumu}}$  measures the amount of error which are accumulated for sequentially running the first T-1 denoising modules.  $\mathbf{E}_t^{\mathrm{cumu}} = D_{\mathrm{KL}} \big( p_{\theta}(\mathbf{x}_{t-1}) \, \| \, q(\mathbf{x}_{t-1}) \big)$  Here  $p_{\theta}(\cdot)$  and  $q(\cdot)$  denote the model distribution in the backward process and reference distribution in the forward process, following the notation of Li & van der Schaar (2024).

**Assumption 2.** (**Predefined Variance Schedule**) As in DDPM or DDIM, we assume each timestep t has a predefined fixed variance  $\sigma_t^2$  that applies uniformly across temporal coordinates. This allows us to simplify the SDE-based error formulation into a constant-variance expression at each step.

**Definition 3.2.1 (Second-order derivative as the numerical error)** We define the *numerical error* (DE) of the first-order derivative solution as

$$\mathbf{x}_t^{\text{Exact}} - \mathbf{x}_t^{\text{Euler(1st)}} \approx \nabla_t \big[ f(t) \, \mathbf{x}_t - \frac{1}{2} g^2(t) \, \nabla_x \log p_t(\mathbf{x}_t) \big].$$

It is the local truncation error of numerically integrating the backward ODE by Euler's method.

**Lemma 3.2.1** (KL gradient in the Gaussian case) For two normal distributions  $\mathcal{P}: \mathcal{N}(\mathbf{x}_t, \sigma_1^2), \mathcal{Q}: \mathcal{N}(\boldsymbol{\mu}_t, \sigma_2^2)$ , this shows that the difference between two Gaussian distributions can be expressed as a distance between their means.

$$\nabla D_{\mathrm{KL}}(\mathcal{P}||\mathcal{Q})) \approx \frac{(\mathbf{x}_t - \boldsymbol{\mu}_t)}{\sigma_2^2}.$$

**Proposition 3.2.1 (Constant Variance Leads to a Simpler Form).** Let  $\sigma^2$  be a predefined constant over the spatial dimensions. Then, the gradient with respect to t of the modified score function from Definition 3.2.1 satisfies

$$\nabla_t \left[ f(t) \mathbf{x}_t - \frac{1}{2} g^2(t) \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t) \right] \approx A \cdot \frac{(\mathbf{x}_t - \boldsymbol{\mu}_t)}{\sigma^2} + B.$$

where  $A, B \in \mathbb{R}$  is the optimization constant to minimize discretization error. With constant variance, the numerical error can be found as a linear equation of the distance of two distributions.

**Theorem 3.2.1 (Approximation Error and KL Gradient Scaling).** For a given step t, the difference between the exact solution  $\mathbf{x}^{\text{Exact}}$  and the Euler discretization  $\mathbf{x}^{\text{Euler}}$  can be expressed as

$$\mathbf{x}^{\mathrm{Exact}} - \mathbf{x}^{\mathrm{Euler}} \approx A \cdot \nabla D_{\mathrm{KL}}(\mathcal{P} \parallel \mathcal{Q}) + B.$$

where  $\mathcal{P}$  is discrete distribution,  $\mathcal{Q}$  is continuous exact distribution. Here, A,B can be derived exactly through linear regression between latant spaces of SR and HR . Intuitively, this tell us that discretization error at each step behaves like the KL gradient between the continuous and discrete processes.

**Discussion.** The detailed proofs of the lemma, proposition, and theorem are provided in the supplemental material. In summary, starting from Assumption 2, we derive Theorem 3.2.1, which establishes that the numerical error—characterized by a second-order derivative—can be formulated as the Kullback–Leibler (KL) divergence between distribution from the discrete process and the continuous process. Assumption 1 defined that cumulative error in a path is defined as divergence of the forward and backward process. Likewise, we assume that exact distribution is substituted with the distribution of the forward process and Euler distribution is with the distribution of the backward process. Theorem 3.2.1 define the error in a single step depends linearly on the gradient of KL divergence of the forward and backward process.

The distribution of the continuous process becomes the reference of the discrete process, which can be substituted with the forward process of HR image because it has better accuracy than the reverse process starting from LR image. Specifically, in the context of a conditional diffusion framework, the backward process is initialized from the low-resolution (LR) image, while the forward process originates from the high-resolution (HR) image. Consequently, we demonstrate that the second-order numerical discretization error exhibits a linear dependency on the KL divergence between the latent representations of the HR and LR images, as governed by the conditional posterior learning model.

In the preceding paragraphs, we defined how the second-order numerical error can be evaluated and corrected. The next step is to formulate the theoretical framework for appyling these equations to the existing posterior model. From Yue et al. (2023), the reverse process for estimating the posterior distribution  $p(\mathbf{x}_0 \mid \mathbf{y}_0)$  is defined by:

$$p(\mathbf{x}_0 \mid \mathbf{y}_0) = \int p(\mathbf{x}_T \mid \mathbf{y}_0) \prod_{t=1}^T p_{\theta}(\mathbf{x}_{t-1} \mid \mathbf{x}_t, \mathbf{y}_0) \, d\mathbf{x}_{1:T}.$$
 (5)

where  $x_0$  is the data distribution and  $y_0$  is a measurement or LR. From above equation, we need to incorporate numerical error correction term from Theorem 3.2.1 into the above equation. For this purpose, we provide an additional LR image as a condition for inference to the existing conditional diffusion model. To achieve this, we extend the ResShift formula with the explicit guided diffusion approach from Dhariwal & Nichol (2021). In Eq. equation 5,  $x_0$  and  $y_0$  represent the HR and LR images, respectively.

$$p(\mathbf{x}_0 \mid \mathbf{y}_0, \mathbf{c}) = \int p(\mathbf{x}_T \mid \mathbf{y}_0, \mathbf{c}) \prod_{t=1}^T p_{\theta}(\mathbf{x}_{t-1} \mid \mathbf{x}_t, \mathbf{y}_0, \mathbf{c}) \, d\mathbf{x}_{1:T},$$
(6)

where c is an additional LR condition.

**Proposition 3.3.1 (Guidance for LR-based Conditioning)** The target distribution  $p(\mathbf{x}_0 \mid \mathbf{y}_0, \mathbf{c})$  with LR guidance is given by

$$\mathcal{N}(\boldsymbol{\mu} + \Sigma \nabla p_{\theta}(\mathbf{c} \mid \mathbf{x}_{t}, \mathbf{y}_{0}, t), \Sigma),$$
 (7)

where c indicates the additional image guidance.

**Remark 3.3.1 (Implementation and Partial Guidance).** From Theorem 3.2.1 and Proposition 3.3.1, one obtains:

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{c} \mid \mathbf{x}_t, \mathbf{y}_0) = \nabla D_{\mathrm{KL}}(p_{\theta}(\mathbf{x}_t) \parallel p_{\theta}(\mathbf{c})), \tag{8}$$

for an image-based guidance approach. This remark connects numerical theory to practice: numerical correction and classifier guidance have the equivalence.

**Discussion.** With Remark 3.3.1, we derive a theoretical basis for using an input image as an additional condition alongside the conventional LR input. In practical scenarios, the conventional condition is passed through a noise-added pipeline as part of the diffusion process, while the newly introduced condition is incorporated through a noise-free path. This dual-conditioning strategy allows us to correct the discretization error accumulated during few-step sampling. Please note that the numerical error correction term is equivalent to our image-based classifier guidance term, up to a constant. This demonstrates that numerical error correction can be realized through image-based classifier guidance, thereby establishing a plug-and-play model. More detailed description about practical usage will be introduced in the next section.

#### 3.4 IMAGE CONTROL VIA IMAGE-BASED GUIDANCE

Based on previous formulations, the second-order numerical error can be represented by the gradient of the KL divergence between the SR and HR distributions. This implies that the second-order numerical error can be linearly regressed using the gradient of the KL divergence, thereby making numerical error restoration theoretically feasible. Simultaneously, from Remark 3.3.1, we defined a method how to supply additional image to the diffusion pipeline as the classifier condition where it is also the gradient of the KL divergence of latant vector and classifier condition. Therefore, if we want to restore numerical error of the posterior model with HR image or LR image, we just need to supply HR image or LR image as classifier condition.

**Reference Region** The following procedure demonstrates how to estimate the constants A and B by measuring the difference in the latent space between SR and HR within the diffusion model and regressing it against the divergence between the two distributions.

$$\Delta = \mu_t(HR) - \mu_t(SR) \tag{9}$$

$$A, B = \arg\min_{a \in \mathbb{R}, b \in \mathbb{R}} \|\Delta - (a \cdot \nabla_x D_{KL} (SR \| HR) + b)\|_2^2$$
(10)

$$\mu_{\text{new}}(SR) = \mu_{\text{old}}(SR) + B + A \cdot \nabla_x D_{\text{KL}}(SR \parallel HR)$$
(11)

where  $A, B \in \mathbb{R}$ .  $\nabla_x D_{\mathrm{KL}}$  (HR || SR<sub>t</sub>) follows the trajectory of the surface between SR and HR.Chen et al. (2024b) Therefore, it can restore truncation error in Theorem 3.2.1. This will be validated in Figure 1, 9, 10.

Non-Reference Region In real-world blind super-resolution, it is not feasible to access HR references, and thus the difference between HR and SR (Eq. 9) as well as the corresponding trajectory (Eq. 10) are unavailable. To address this, we propose to estimate the parameters A and B from regions where LR–HR pairs exist, and then apply them to LR-only regions. Since Eq. 10 does not exist in such settings, we instead utilize  $\nabla_x D_{\rm KL}({\rm SR}|{\rm LR})$ . While Eq. 10 corresponds to the gradient at each time step along the SR–HR trajectory, if the LR–SR trajectory resided in the extrapolated space of the SR–HR trajectory, one could substitute  $D_{\rm KL}({\rm SR}|{\rm HR})$  with  $D_{\rm KL}({\rm SR}|{\rm LR})$  and compute A and B to construct the correction term. Our validation experiments in Fig. 5 confirmed that they are well aligned in the same trajectory, particularly in the pretrained model of Yue et al. (2023). Instead, we propose a scheme that approximates numerical errors using  $D_{\rm KL}({\rm SR}|{\rm LR})$ . This scheme is summarized in Algorithm 1 and Algorithm 2.

The process of determining parameters A and B can be regarded as a calibration step, where multiple LR–HR image pairs are used to identify numerical errors and extract the corresponding parameters, which are then provided during inference. Linear regression over multiple images failed to yield optimized parameters for both A and B. Consequently, we abandon the estimation of B and instead focus solely on finding A, which minimizes numerical errors across multiple images and is supplied during inference. More details about how to derive algorithm 1 and algorithm 2 will be provided in Fig. 4 and in section 6.

# 4 EXPERIMENT

**Testing Environment** To validate our theoretical framework, we conducted experiments using the DIV2K validation dataset Agustsson & Timofte (2017), the RealSR dataset Cai et al. (2019b), and Flickr30k Dataset Young et al. (2014). Since datasets contain images of varying sizes, we cropped  $128 \times 128$  center region and further divided them into four  $128 \times 128$  patches. We used 100 images for DIV2K and RealSR for generation of quantitative table. Calibration of Guidance scale is performed over Imagenet Dataset Deng et al. (2009) which is used in training the pretrained model of Yue et al. (2023). We conducted all experiments on Nvidia H100 machine.

**Metrics for comparison** For a comprehensive image quality assessment, we employ both full-reference and no-reference metrics. Reference-based fidelity measures: PSNR and SSIM Wang et al. (2004) / Reference-based perceptual quality measures: LPIPS Zhang et al. (2018) and DISTS Ding et al. (2020) / No-reference image quality measures: NIQE Zhang et al. (2015), MANIQA Yang et al. (2022), and CLIPIQA Wang et al. (2023b) / Image distribution-based metric: FID Heusel et al. (2017), which evaluates the distance between restored images and ground truth.

#### Algorithm 1 Calibrate Guidance Scale

378

379

380

381

382

384

386

387

388 389

390

391

392 393

394

406

407

408

409

410

411

412 413

414

415

416

417

418

419

420

421

422

423

424

425

426

427

428

429

430

431

```
Require: A set of LR–HR pairs (\mathbf{LR}_i, \mathbf{HR}_i) model
       parameters \theta, candidate scale set \{\alpha_t\}
Ensure: A chosen guidance scale \alpha_t^*.
 1: for i = 1 to n do
2: \mathcal{Z}_{T}^{(\mathrm{HR})} \leftarrow \mathrm{ForwardODE}(\mathbf{HR}_{i})
 3:
            x_T \leftarrow \mathbf{L}\mathbf{R}_i
            for t = T down to 1 do
  5:
                 (\mu, \Sigma) \leftarrow \mu_{\theta}(x_t), \ \Sigma_{\theta}(x_t)
                 \Delta = \mu_{\theta}(\mathcal{Z}_{T}^{(HR)}) - \mu
                G_{\mathrm{LR}} = \nabla_{x_t} D_{\mathrm{KL}} (p_{\theta}(x_t) \parallel \mathbf{LR}_i)
                 L_i(\alpha_t) = \|\Delta\| - \|\alpha_t \cdot G_{LR}\|
 9:
            end for
10: end for
11: \alpha_t^* = \underset{\alpha_t}{\operatorname{argmin}} \sum_{t=1}^n L_i(\alpha_t).
```

#### **Algorithm 2** Inference with LR Guidance

```
Require: LR image (or condition) c, guidance scale \alpha_t^*, diffusion model parameters \theta,

Ensure: Reconstructed SR image x_0.

1: x_T \leftarrow \mathbf{LR}_i
2: for t = T down to 1 do
3: (\mu, \Sigma) \leftarrow \mu_{\theta}(x_t), \Sigma_{\theta}(x_t)
4: x_{t-1} \leftarrow
5: \mathcal{N}(\mu + \alpha_t^* \Sigma \nabla_{x_t} D_{\mathrm{KL}}(p_{\theta}(x_t) \parallel \mathbf{LR}_i))
6: end for
7: return x_0
```

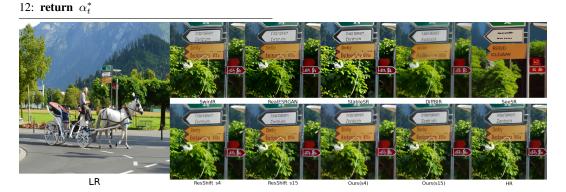


Figure 1: **Reference Region**. Ours (s4) and Ours (s15) represent ResShift (s4 and s15) with the proposed correction module in Eq. 11. In the region with reference overlap, the KL divergence gradient between HR and SR is used to restore the numerical error. This demonstrates that the proposed method effectively corrects the error in the overlapped region. After numerical error correction, architectural features such as the text on travel signboards were successfully restored. In particular, the s15 model achieved restoration results nearly indistinguishable from the HR image because its truncation error is already lower than that of the s4 model.

Quantitative Comparison with State-of-the-art methods To validate the effectiveness of our proposed method, we conducted quantitative and qualitative comparisons against various state-ofthe-art (SOTA) methods, including: SwinIR Liang et al. (2021b), Real-ESRGAN Wang et al. (2021), StableSR Wang et al. (2023c), ResShift Yue et al. (2023), PASD Yang et al. (2024), DiffBIR Lin et al. (2024), SinSR Wang et al. (2024b), OseDiff Wu et al. (2024a), InvSR Yue et al. (2025). Our comparative analysis includes regression-based, GAN-based, and diffusion-based super-resolution methods. In Table 2, the best and second-best results are highlighted in red and blue, respectively, while regression- and GAN-based methods were excluded from this ranking. To the best of our knowledge, ResShift Yue et al. (2023) is currently the only posterior-based super-resolutin model. Thereby, we adopt ResShift Yue et al. (2023) as our backbone framework. Other methods based on pretrained diffusion prior such as Rombach et al. (2022b) are PASD Yang et al. (2024), StableSR Wang et al. (2023c), and DiffBIR Lin et al. (2024). They incorporate additional neural networks into the diffusion prior model. This could lead to overlapping effects between our theoretical enhancement and their inherent framework improvements. From Table 2, our algorithm achieves SOTA performance in non-reference metrics such as MANIQA, CLIPIQA, and FID, as well as in perception-related reference metrics such as LPIPS while minimizing degradation in reference based metrics. This demonstrates that our method effectively restores the perceptual quality of posterior models. Our method is implemented as a plug-in module integrated into the ResShift-4s or ResShift-15s framework. As such, its performance fundamentally converges to the quality of the base model. Due to discretization error, the image quality of ResShift-4s is inherently lower than that of the ResShift-15s model and can be identified in Figure 2. In the HR ground truth, all wires



Figure 2: **Non-Reference Region** Qualitative comparisons in the non-reference region. Ours shows perceptual quality enhancement in the non-overlapped region with the reference according to the Remark 3.3.1

Table 2: Quantitative comparison with contemporary state-of-the-art image super resolution algorithms for non-overlapped region. 's' denotes the number of diffusion steps. The best and second-best results of each metric are highlighted in red and blue, respectively.

Datasets	Methods	PSNR↑	SSIM↑	LPIPS↓	DISTS↓	NIQE↓	MANIQA↑	CLIPIQA↑	FID↓
RealSR	SwinIR	26.401153	0.822698	0.198882	0.212861	5.533595	0.482269	0.520561	34.319923
	Realesrgan	24.824283	0.690881	0.228614	0.211228	5.554380	0.515528	0.602108	68.559638
	StableSR(s200)	23.141011	0.700490	0.251224	0.218401	5.993226	0.445875	0.619377	66.324215
	DiffBIR(s50)	24.044400	0.680713	0.274067	0.221713	5.538441	0.618653	0.707664	61.714618
	PASD(s20)	24.608645	0.712806	0.242913	0.209351	4.795048	0.540107	0.625039	52.826575
	Resshift(s15)	24.932412	0.734073	0.198175	0.187559	5.319726	0.557401	0.677602	46.756755
	Resshift(s4)	25.262886	0.729898	0.207532	0.193287	5.930618	0.481287	0.646861	47.521572
	SinSR(s1)	24.814703	0.732119	0.264237	0.204537	5.189266	0.504561	0.639777	49.655490
	Osediff(s1)	23.221523	0.658442	0.279372	0.221037	5.464938	0.477634	0.661837	69.084736
	InvSR(s1)	23.163527	0.660025	0.274930	0.226382	4.827873	0.588372	0.704826	45.806334
	Ours(s4, HR)	27.563256	0.859821	0.169382	0.232162	5.158372	0.553831	0.651693	35.052837
	Ours(s4, LR)	24.537216	0.730573	0.201882	0.188932	5.366422	0.649319	0.757744	44.208402
DIV2k	SwinIR	25.524262	0.756364	0.196358	0.253459	6.574358	0.236129	0.369175	66.022553
	Realesrgan	25.017905	0.729274	0.208930	0.231028	4.377671	0.323017	0.440141	64.281207
	StableSR(s200)	25.089883	0.723934	0.230532	0.241219	6.143113	0.334181	0.402819	63.208495
	DiffBIR(s50)	24.952325	0.664533	0.245898	0.222816	5.170970	0.519603	0.694852	64.320329
	PASD(s20)	25.268907	0.730030	0.239548	0.222048	5.840237	0.409314	0.503350	66.624089
	Resshift(s15)	25.195770	0.682318	0.236521	0.230849	6.443043	0.367619	0.565044	69.377387
	Resshift(s4)	25.195317	0.682465	0.210933	0.236403	6.325634	0.324337	0.526237	68.193004
	SinSR(s1)	25.323979	0.700209	0.226813	0.238258	6.067024	0.359430	0.557111	87.956015
	Osediff(s1)	24.794832	0.714938	0.2194821	0.213805	5.859347	0.427836	0.689721	67.801938
	InvSR(s1)	24.143282	0.670584	0.2415381	0.253827	5.981039	0.471184	0.691187	74.120972
	Ours(s4,LR)	24.910596	0.682398	0.215568	0.235888	6.024669	0.475017	0.707421	68.141436

are structured as dual wires, but among the SR results, only SwinIR, ResShift-15s, and our method successfully reconstruct the dual-wire structure. Naturally, a 15-step model is expected to suffer less from discretization error than a 4-step model (as long as the model is properly trained). Nonetheless, despite operating with only four steps, our method succeeds in reconstructing the dual-wire structure by effectively restoring its discretization error. Our framework is the first to provide a controllable knob for navigating this trade-off without retraining.

Ablation Study, Qualitative Results is available in section 6 and 8 in the supplemental material.

#### 5 Conclusion

In this work, we introduced a plug-and-play module for posterior learning-based super-resolution models. From a theoretical perspective, we derived that discretization errors inherent in first-order ODE formulations can be corrected by incorporating second-order differential terms. We further showed that this correction can be formulated within the conditional diffusion framework, and that the resulting expression naturally coincides with the structure of image-based classifier guidance. This provides a new theoretical understanding of posterior diffusion trajectories.

From a practical perspective, we demonstrated that applying the correction term can reduce numerical errors in pretrained models without retraining. While fidelity enhancement through this correction requires HR supervision and is thus limited in blind SR scenarios, we proposed instead to exploit LR-based classifier guidance in a sign-flipped manner to enhance perceptual quality. This design enables posterior SR models to achieve perceptual results competitive with state-of-the-art methods, thereby bridging theoretical insight with practical applicability.

# REFERENCES

- Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *CVPRW*, July 2017.
  - Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Deep burst super-resolution. In CVPR, 2021.
  - Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *ICCV*, pp. 3086–3095, 2019a.
  - Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *ICCV*, pp. 3086–3095, 2019b.
    - Bin Chen, Gehui Li, Rongyuan Wu, Xindong Zhang, Jie Chen, Jian Zhang, and Lei Zhang. Adversarial diffusion compression (adc) for real-world isr. In *CVPR*, 2025.
  - Defang Chen, Zhenyu Zhou, Jian-Ping Mei, Chunhua Shen, Chun Chen, and Can Wang. A geometric perspective on diffusion models. In *ICLR*, May 2024a.
  - Defang Chen, Zhenyu Zhou, Can Wang, Chunhua Shen, and Siwei Lyu. On the trajectory regularity of ode-based diffusion sampling. In *ICML*, July 2024b.
  - Junyang Chen, Jinshan Pan, and Jiangxin Dong. Faithdiff: Faithful and controllable image super-resolution with diffusion models. 2024c.
  - Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In CVPR, pp. 248–255, 2009.
    - Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, volume 34, pp. 8780–8794, 2021.
    - Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying structure and texture similarity. In *IEEE TPAMI*, volume 44, pp. 2567–2581, 2020.
    - Linwei Dong, Qingnan Fan, Yihong Guo, Zhonghao Wang, Qi Zhang, Jinwei Chen, Yawei Luo, and Changqing Zou. Tsd-sr: One-step diffusion with target score distillation for real-world image super-resolution. 2024.
    - Weisheng Dong, Lei Zhang, Guangming Shi, and Xiaolin Wu. Image deblurring and super-resolution by adaptive sparse domain selection and adaptive regularization. *IEEE TIP*, 20(7):1838–1857, 2011.
    - Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, pp. 12873–12883, 2021.
    - Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. In *NeurIPS*, volume 2, pp. 2672–2680, 2014.
    - Shuhang Gu, Wangmeng Zuo, Qi Xie, Deyu Meng, Xiangchu Feng, and Lei Zhang. Convolutional sparse coding for image super-resolution. In *ICCV*, pp. 1823–1831, 2015.
    - Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
    - Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, volume 33, pp. 6840–6851, 2020.
    - M. Hochbruck and A. Ostermann. Explicit exponential runge-kutta methods for semilinear parabolic problems. In SIAM Journal on Numerical Analysis, volume 43, pp. 1069–1090, 2005.
    - et al. Joram Soch. Statproofbook/statproofbook.github.io: The book of statistical proofs. 2024.
    - Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Proc. NeurIPS*, 2022.
- Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In
   NeurIPS, 2022.
  - Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In CVPR, 2017.

- Junyong Lee, Myeonghee Lee, Sunghyun Cho, and Seungyong Lee. Reference-based video super-resolution using multi-camera video triplets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
  - Haoying Li, Yifan Yang, Meng Chang, Shiqi Chen, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing*, 479:47–59, 2022a.
  - Haoying Li, Yifan Yang, Meng Chang, Shiqi Chen, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing*, 479:47–59, 2022b.
  - Haoying Li, Yifan Yang, Meng Chang, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. Srdiff: Single image super-resolution with diffusion probabilistic models. volume 479, pp. 47–59, 2022c.
  - Mingxiao Li, Tingyu Qu, Ruicong Yao, Wei Sun, and Marie-Francine Moens. Alleviating exposure bias in diffusion models through sampling with shifted time steps. In *ICLR*, May 2024.
  - Yangming Li and Mihaela van der Schaar. On error propagation of diffusion models. In ICLR, May 2024.
  - Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *ICCV-W*, pp. 1833–1844, 2021a.
  - Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *ICCV-W*, pp. 1833–1844, 2021b.
  - Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPR*, 2017.
  - Xinqi Lin, Jingwen He, Ziyan Chen, Zhaoyang Lyu, Bo Dai, Fanghua Yu, Wanli Ouyang, Yu Qiao, and Chao Dong. Diffbir: Towards blind image restoration with generative diffusion prior, 2024.
  - Jiawei Liu, Qiang Wang, Huijie Fan, Yinong Wang, Yandong Tang, and Liangqiong Qu. Residual denoising diffusion models. In *CVPR*, 2024.
  - Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. In *ICLR*, 2022.
  - Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 5 steps. In *NeurIPS*, 2022.
  - Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. In *ICLR*, 2023.
  - Nancy Mehta, Akshay Dudhane, Subrahmanyam Murala, Syed Waqas Zamir, Salman Khan, and Fahad Shahbaz Khan. Gated multi-resolution transfer network for burst restoration and enhancement. In *CVPR*, 2023.
  - E. Kloeden Peter and Platen Eckhard. Numerical Solution of Stochastic Differential Equations. Springer, 1992.
  - Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *CVPR*, 2021.
  - Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10684–10695, 2022a.
  - Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10684–10695, 2022b.
  - Chitwan Saharia, Jonathan Ho, William Chan, David J. Fleet, Mohammad Norouzi, and Tim Salimans. Image super-resolution via iterative refinement. In *CVPR*, 2021.
  - Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH Conference*, pp. 1–10, 2022a.
    - Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE TPAMI*, 2022b.
    - Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE TPAMI*, 2022c.

- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021a.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021b.
  - Gilbert Strang. Computational Science and Engineering. Wellesley-Cambridge Press, 2007.
  - Hang Wang, Xuanhong Chen, Bingbing Ni, Yutian Liu, and Jinfan Liu. Omni aggregation networks for lightweight image super-resolution. In CVPR, 2023a.
  - Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023b.
  - Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin CK Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. *arXiv preprint arXiv:2305.07015*, 2023c.
  - Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *ECCV-W*, 2018.
  - Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *ICCV-W*, pp. 1905–1914, 2021.
  - Yufei Wang, Wenhan Yang, Xinyuan Chen, Yaohui Wang, Lanqing Guo, Lap-Pui Chau, Ziwei Liu, Yu Qiao, Alex C. Kot, and Bihan Wen. Sinsr: Diffusion-based image super-resolution in a single step. In CVPR, 2024a.
  - Yufei Wang, Wenhan Yang, Xinyuan Chen, Yaohui Wang, Lanqing Guo, Lap-Pui Chau, Ziwei Liu, Yu Qiao, Alex C Kot, and Bihan Wen. Sinsr: diffusion-based image super-resolution in a single step. In *CVPR*, pp. 25796–25805, 2024b.
  - Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004.
  - Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. 2024a.
  - Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. Seesr: Towards semantics-aware real-world image super-resolution. In *CVPR*, pp. 25456–25467, 2024b.
  - Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assessment. In *CVPR*, pp. 1191–1200, 2022.
  - Tao Yang, Rongyuan Wu, Peiran Ren, Xuansong Xie, and Lei Zhang. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. In *ECCV*, 2024.
  - Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. In *TACL*, volume 2, pp. 67–78, 2014.
  - Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Resshift: Efficient diffusion model for image super-resolution by residual shifting. In *NeurIPS*, 2023.
  - Zongsheng Yue, Kang Liao, and Chen Change Loy. Arbitrary-steps image super-resolution via diffusion inversion. In *CVPR*, 2025.
  - Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, 2022.
  - Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *ICCV*, pp. 4791–4800, 2021.
  - Lin Zhang, Lei Zhang, and Alan C Bovik. A feature-enriched completely blind image quality evaluator. In *IEEE TIP*, volume 24, pp. 2579–2591, 2015.
  - Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pp. 586–595, 2018.
  - Kaiwen Zheng, Cheng Lu, Jianfei Chen, and Jun Zhu. Dpm-solver-v3: Improved diffusion ode solver with empirical model statistics. In NeurIPS, 2023.