

# TSD-SR: One-Step Diffusion with Target Score Distillation for Real-World Image Super-Resolution

Linwei Dong<sup>1,2\*</sup> Qingnan Fan<sup>2\*</sup> Yihong Guo<sup>1</sup> Zhonghao Wang<sup>3</sup>  
Qi Zhang<sup>2</sup> Jinwei Chen<sup>2</sup> Yawei Luo<sup>1†</sup> Changqing Zou<sup>1,4</sup>

<sup>1</sup>Zhejiang University <sup>2</sup>Vivo Mobile Communication Co. Ltd  
<sup>3</sup>University of Chinese Academy of Sciences <sup>4</sup>Zhejiang Lab

## Abstract

Pre-trained text-to-image diffusion models are increasingly applied to real-world image super-resolution (Real-ISR) task. Given the iterative refinement nature of diffusion models, most existing approaches are computationally expensive. While methods such as SinSR and OSEDiff have emerged to condense inference steps via distillation, their performance in image restoration or details recovery is not satisfied. To address this, we propose TSD-SR, a novel distillation framework specifically designed for real-world image super-resolution, aiming to construct an efficient and effective one-step model. We first introduce the Target Score Distillation, which leverages the priors of diffusion models and real image references to achieve more realistic image restoration. Secondly, we propose a Distribution-Aware Sampling Module to make detail-oriented gradients more readily accessible, addressing the challenge of recovering fine details. Extensive experiments demonstrate that our TSD-SR has superior restoration results (most of the metrics perform the best) and the fastest inference speed (e.g. 40 times faster than SeeSR) compared to the past Real-ISR approaches based on pre-trained diffusion priors.

## 1. Introduction

Image super-resolution (ISR) [8, 9, 23, 26] aims to transform low-quality (LQ) images, which have experienced noise or blur, into clear high-quality (HQ) images. Differing from traditional ISR [6, 66], which assumes a known degradation process, real-world image super-resolution (Real-ISR) [46, 62] is designed to enhance real-world images that have suffered from complex and unknown degradations, thereby offering greater practical utility.

Generative models, particularly Generative Adversarial Networks (GANs) [11, 33, 37] and Diffusion Models (DMs)

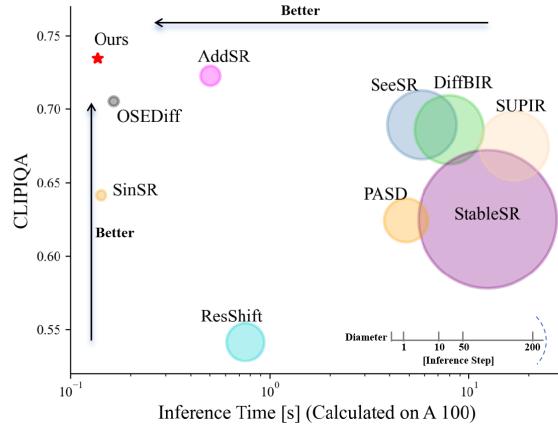


Figure 1. Performance and efficiency comparison among Real-ISR methods. TSD-SR stands out for achieving high-quality restoration with the fastest speed among diffusion-based models. In contrast, existing models prioritize either speed or restoration performance. The performance of each method is benchmarked on an A100 GPU with the DRealSR dataset.

[17, 39, 41], have shown remarkable power in handling Real-ISR task. GAN-based methods leverage adversarial training, toggling between the generator and discriminator to produce realistic images. While GANs are capable of one-step inference, they are commonly hampered by issues like mode collapse and training instability [2]. Recently, Diffusion Models (DMs) have shown impressive performance in the realm of image generation [21, 47]. Their robust priors empower them to produce more realistic images with richer details than GAN-based methods [39, 41]. Some researchers [29, 53, 56, 60] have successfully leveraged pre-trained DMs for Real-ISR task. However, due to the iterative denoising nature of diffusion models [17], the Real-ISR process is computationally expensive.

To achieve an efficient and one-step network akin to GANs, several pioneering methods that condense the iter-

\*Equal contribution. †Corresponding authors.

ations of diffusion models through distillation [12, 15, 18, 57] have been proposed [48, 52, 54]. Among these works, OSEDiff [52] introduced the Variational Score Distillation (VSD) loss [50] to Real-ISR task, achieving state-of-the-art (SOTA) one-step performance by leveraging prior knowledge from pre-trained models. Despite these advancements, our investigation has revealed two critical limitations associated with VSD in Real-ISR applications. **(1) Unreliable gradient direction.** VSD relies on a Teacher Model to provide a “true gradient direction.” However, this guidance proves unreliable in scenarios where initial ISR outputs are suboptimal. **(2) Insufficient detail recovery.** The VSD loss exhibits notable variation across different timesteps, and the uniform sampling strategy for  $t$  poses challenges in aligning the score function with detailed texture recovery requirements. These findings highlight the need for enhanced approaches to address these issues effectively.

In this paper, we propose a novel method called **TSD-SR** to distill a multi-step Text-to-Image (T2I) DMs [10, 38, 39] into an effective one-step diffusion model tailored for the Real-ISR task. Specifically, TSD-SR consists of two components: **Target Score Distillation (TSD)** and **Distribution-Aware Sampling Module (DASM)**. TSD incorporates our new proposed Target Score Matching (TSM) loss to compensate for the limitation of VSD loss. This significant score loss leverages HQ data to provide a reliable optimization trajectory for the distillation process. It effectively reduces the visual artifacts caused by deviant predictions from the Teacher Model. DASM is crafted to bolster detail recovery by strategically sampling low-noise samples that are distribution-based during training. This approach effectively allocates more optimization to early timesteps in a single iteration, enhancing the recovery of details.

Experiments on popular benchmarks demonstrate that TSD-SR achieves superior restoration performance (most of the metrics perform the best) and high efficiency (the fastest inference speed, 40 times faster than SeeSR) compared to the state-of-the-art Real-ISR methods based on pre-trained DMs, while requiring only a single inference step.

Our main contribution can be summarized as threefold:

- We propose a novel method called TSD-SR to achieve one-step DMs distillation for the Real-ISR task.
- We introduce Target Score Distillation (TSD) to provide a reliable gradient to enhance the realism of the Real-ISR model’s outputs.
- We design a Distribution-Aware Sampling Module (DASM) specifically tailored to enhance the capability of detail restoration.

## 2. Related Work

**GAN-based Real-ISR.** Since SRGAN [26] first applied GAN to ISR, it has effectively enhanced visual quality by combining adversarial loss with perceptual loss [7,

65]. Subsequently, ESRGAN [45] introduced Residual-in-Residual Dense Block and a relativistic average discriminator, further improving detail restoration. Methods like BSRGAN [62] and Real-ESRGAN [46] simulate complex real-world degradation processes, achieving ISR under unknown degradation conditions, which enhances the model’s generalization ability. Although GAN-based methods are capable of adding more realistic details to images, they suffer from training instability and mode collapse [2].

**Multi-step Diffusion-based Real-ISR.** Some researches [29, 44, 53, 56, 60] in recent years have utilized the powerful image priors in pre-trained T2I diffusion models [35, 39, 64] for Real-SR tasks and achieved promising results. For example, StableSR [44] balances fidelity and perceptual quality by fine-tuning the time-aware encoder and employing controllable feature wrapping. DiffBiR [29] first processes the LR image through a reconstruction network and then uses the Stable Diffusion (SD) model [39] to supplement the details. SeeSR [53] attempts to better stimulate the generative power of the SD model by extracting the semantic information in the image as a conditional guide. PASD [56] introduces a pixel-aware cross attention module to enable the diffusion model to perceive the local structure of the image at the pixel level, while using a degradation removal module to extract degradation insensitive features to guide the diffusion process along with high-level information from the image. SUPIR [60] achieves a generative and fidelity capability using negative cues [16] as well as restoration-guided sampling, while using a larger pre-training model with a larger dataset to enhance the model capability. However, all of these methods are limited by the multi-step denoising of the diffusion model, which requires 20-50 iterations in inference, resulting in an inference time that lags far behind that of GAN-based methods.

**One-step Diffusion-based Real-ISR.** Recently, there has been a surge of interest within the academic community in one-step distillation techniques [31, 34, 40, 58, 59] for diffusion-based Real-ISR task. SinSR [48] leverages consistency preserving distillation to condense the inference steps of ResShift [61] into a single step, yet the generalization of ResShift and SinSR is constrained due to the absence of large-scale data training. AddSR [54] introduces the adversarial diffusion distillation (ADD) [40] to Real-ISR tasks, resulting in a comparatively effective four-step model. However, this method has a propensity to produce excessive and unnatural image details. OSEDiff [52] directly uses LQ images as the beginning of the diffusion process, and employs VSD loss [50] as a regularization technique to condense a multi-step pre-trained T2I model into a one-step Real-ISR model. However, due to the incorporation of alternating training strategies, OSEDiff may initially tend towards unreliable optimization directions, which may lead to the visual artifacts.

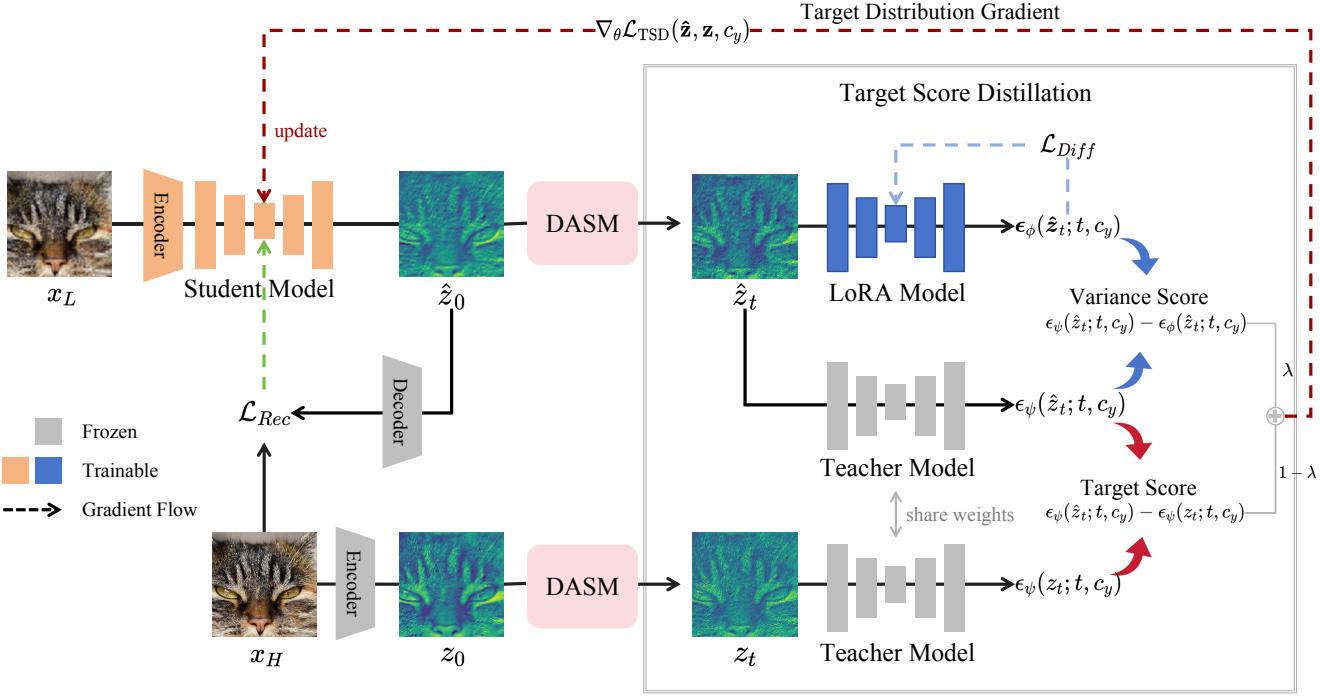


Figure 2. Pipeline overview. We train one-step Student Model  $G_\theta$  to map the low-quality image  $x_L$  into a more realistic one. The noisy latent  $\hat{z}_t$  and  $z_t$  sampled by DASM (Details can be found in Fig. 6.) will be fed into both the pre-trained Teacher and LoRA Model to produce the Variational Score loss. Subsequently, two Models’ predictions on  $z_t$  yield Target Score loss. Their weighted forms, namely TSD (red flow), in conjunction with pixel-space reconstruction loss (green flow), are leveraged to update the Student Model. After updating the Student Model, we employ the diffusion loss (blue flow) to alternately update the LoRA Model.

### 3. Methodology

#### 3.1. Preliminaries

**Problem Formulation.** ISR problem aims to reconstruct a HQ image  $x_H$  from a LQ input  $x_L$  by training a parameterized ISR model  $G_\theta$  on a dataset  $\mathcal{D} = \{(x_L, x_H)_{i=1}^N\}$ , where  $N$  represents the number of image pairs. It is mathematically formulated as minimizing the following objective:

$$\theta^* = \arg \min_{\theta} \mathbb{E}_{(x_L, x_H) \sim \mathcal{D}} [\mathcal{L}_{Rec}(G_\theta(x_L), x_H) + \lambda \mathcal{L}_{Reg}(q_\theta(\hat{x}_H), p(x_H))] \quad (1)$$

Here,  $\mathcal{L}_{Rec}$  denotes the reconstruction loss, commonly measured using distance metrics like  $L_2$  or  $LPIPS$  [65]. The regularization term  $\mathcal{L}_{Reg}$  enhances the realism and generalization of the ISR model’s outputs. This objective can be understood as aligning the ISR output  $\hat{x}_H$ ’s distribution,  $q_\theta(\hat{x}_H)$ , with the true distribution  $p(x_H)$  by minimizing the KL-divergence [25]:

$$\min_{\theta} \mathcal{D}_{KL}(q_\theta(\hat{x}_H) \| p(x_H)) \quad (2)$$

While several studies [45, 46, 62] have employed adversarial loss to optimize this objective, they often encounter issues like mode collapse and training instability. Recent work [52] achieved the state-of-the-art results using Variational Score Distillation (VSD) as the regularization loss to minimizing this objective, which inspires our research.

**Variational Score Distillation.** Variational Score Distillation (VSD) [50] was initially introduced for text-to-3D generation by distilling a pre-trained text-to-image diffusion model to optimize a single 3D representation [36].

In the VSD framework, a pre-trained diffusion model, represented as  $\epsilon_\psi$ , and its trainable (LoRA [19]) replica  $\epsilon_\phi$ , are used to regularize the generator network  $G_\theta$ . As outlined in ProlificDreamer [50], the gradient with respect to the generator parameters  $\theta$  is formulated as follows:

$$\begin{aligned} & \nabla_{\theta} \mathcal{L}_{VSD}(\hat{z}, c_y) \\ &= \mathbb{E}_{t, \epsilon} \left[ w(t) (\epsilon_\psi(\hat{z}_t; t, c_y) - \epsilon_\phi(\hat{z}_t; t, c_y)) \frac{\partial \hat{z}}{\partial \theta} \right] \quad (3) \end{aligned}$$

where  $\hat{z}_t = \alpha_t \hat{z} + \sigma_t \epsilon$  is the noisy input,  $\hat{z}$  is the latent outputted by the generator network  $G_\theta$ ,  $\epsilon$  is a gaussian noise, and  $\alpha_t, \sigma_t$  are the noise-data scaling constants.  $c_y$  is a text embedding corresponding to a caption that describes the input image, and  $w(t)$  is a time-varying weighting function.

#### 3.2. Overview of TSD-SR

As depicted in Fig. 2, our goal is to distill a given pre-trained T2I DMs into a fast one-step Student Model  $G_\theta$  with Teacher Model  $\epsilon_\psi$  and trainable LoRA Model  $\epsilon_\phi$ . We denote the latent outputs of the distilled model as  $\hat{z}_0$ , and HQ latent representations as  $z_0$ . Both  $\hat{z}_0$  and  $z_0$  will pass through our Distribution-Aware Sampling Module (DASM)

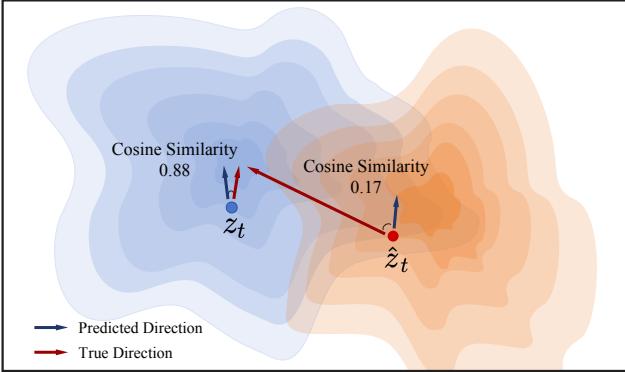


Figure 3. A visual comparison of the gradient direction. We set the timestep  $t$  to 100 and calculated the cosine similarity between the predicted directions from the Teacher Model and the true direction (towards the HQ data). The **prediction direction** for  $z_t$  closely matches the **true direction**, but not for  $\hat{z}_t$ , suggesting that suboptimal samples may lead to directional deviations.

to obtain distribution-based samples  $\hat{z}_t$  and  $z_t$  (Sec. 3.4). We train  $G_\theta$  by minimizing the two losses: a reconstruction loss in pixel space to compare the model outputs against the ground truth, and a regularization loss (from Target Score Distillation) to enhance the realism (Sec. 3.3). After updating the Student Model, we update the LoRA Model with diffusion loss. Finally, in Sec. 3.5, we present an overview of all the losses encountered during the training phase.

### 3.3. Target Score Distillation

Similar to [52], we introduce VSD loss into our work as a regularization loss to enhance the realism and generalization of the  $G_\theta$  outputs. Upon reviewing VSD Eq. (3),  $\epsilon_\phi(\hat{z}_t; t, c_y)$  represents the gradient direction of the optimization on  $G_\theta$  noisy outputs  $\hat{z}_t$ , whereas  $\epsilon_\psi(\hat{z}_t; t, c_y)$  corresponds to the gradient direction on noisy HQ latent representations  $z_t$ . The overarching goal of model optimization is to harmonize the predictions from  $\hat{z}_t$  with those from  $z_t$ . Nonetheless, this strategy encounters hurdles, especially in the early training phase: the quality of synthetic latents  $\hat{z}_t$  is not high enough for Teacher Model to give a precise prediction. As illustrated in Fig. 3, the Teacher Model finds it challenging to accurately predict the optimization direction for low-quality synthetic latents, with the cosine similarity to the ideal optimization direction only 0.2 (0.88 on HQ latents). This problem can lead to severe visual artifacts, as evident in Fig. 4(a).

A straightforward remedial measure is to employ a mean squared error (MSE) loss to align the synthetic latents with the ideal inputs of the Teacher Model, which are derived from HQ latents. However, as shown in Fig. 4(b), this approach has been observed to lead to over-smoothed results [13]. Our strategy, instead, is to align the predictions made by the Teacher Model on both synthetic and HQ

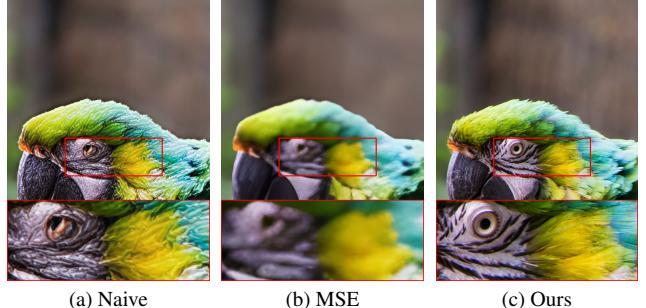


Figure 4. The visualization of different strategies. (a) The naive method introduces fake textures and fails to recover fine details. (b) MSE leads to over-smoothed generation result, lacking high-frequency information. (c) Our method offers the superior visual effects and fine textures.

latents, thereby encouraging greater consistency between them. The core idea is that for samples drawn from the same distribution, the real scores predicted by the Teacher Model should be close to each other. We refer to this approach as Target Score Matching (TSM):

$$\begin{aligned} & \nabla_\theta \mathcal{L}_{\text{TSM}}(\hat{\mathbf{z}}, \mathbf{z}, c_y) \\ &= \mathbb{E}_{t, \epsilon} \left[ w(t) (\epsilon_\psi(\hat{\mathbf{z}}_t; t, c_y) - \epsilon_\psi(\mathbf{z}_t; t, c_y)) \frac{\partial \hat{\mathbf{z}}}{\partial \theta} \right] \end{aligned} \quad (4)$$

where the expectation of the gradient is computed across all diffusion timesteps  $t \in \{1, \dots, T\}$  and  $\epsilon \sim \mathcal{N}(0, I)$ . Equation (4) encapsulates the optimization loss for our Target Score Matching. Upon examining it in conjunction with the Eq. (3), we notice that VSD utilizes the prediction residual between a Teacher and a LoRA Model to drive gradient backpropagation. Conversely, our TSM employs the synthetic and the HQ data to produce the gradients. By blending these two strategies with hyperparameter weighting  $\lambda$  and  $1 - \lambda$ , we can construct our ultimate optimization loss, effectively unifying the strengths of both approaches to optimize the training process Eq. (5).

$$\nabla_\theta \mathcal{L}_{\text{TSD}}(\hat{\mathbf{z}}, \mathbf{z}, c_y) = \mathbb{E}_{t, \epsilon} \left[ w(t) [\epsilon_\psi(\hat{\mathbf{z}}_t; t, c_y) - \epsilon_\psi(\mathbf{z}_t; t, c_y) + \lambda (\epsilon_\psi(\mathbf{z}_t; t, c_y) - \epsilon_\phi(\hat{\mathbf{z}}_t; t, c_y))] \frac{\partial \hat{\mathbf{z}}}{\partial \theta} \right] \quad (5)$$

where  $w(t)$  is time-aware weighting function tailored for Real-ISR. Other symbols are in accordance with those previously mentioned. By introducing the prediction of pre-trained diffusion model on HQ latents, we have circumvented the issue of the model falling into the visual artifacts or over-smoothed problem, as illustrated in Fig. 4(c).

### 3.4. Distribution-Aware Sampling Module

In the VSD-based framework, we need to match the score functions predicted by the Teacher Model with those by the LoRA Model at different timesteps  $t \sim \{0, 1, \dots, T\}$ .

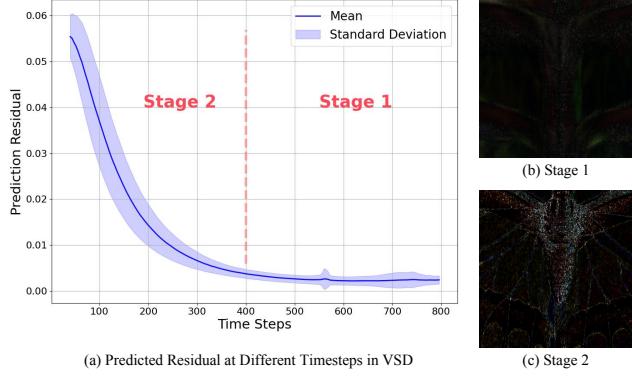


Figure 5. (a) The prediction errors of the VSD loss at different timesteps. The error divergence is more pronounced in early timesteps than later. This phenomenon is observed throughout the optimization process. (b) the visualization of Stage 1 prediction error. (c) the visualization of Stage 2 prediction error.

However, for the Real-ISR problem, this matching ability at different  $t$  is unbalanced, as illustrated in Fig. 5(a). This phenomenon could be attributed to our leveraging of the low-frequency (LF) priors from LQ data, but losing guidance from high-frequency (HF) details. The noisy samples  $\hat{z}_t$  possesses low-frequency priors (from LQ data). These priors are easily learned by the LoRA Model, leading to more similar predictions (Fig. 5(b)) during LF restoration (Stage 1). However, In the Stage 2, the lack of HF details in  $\hat{z}_t$  causes the LoRA Model to poorly perceive details, resulting in divergent predictions (Fig. 5(c)). Our goal is to reduce this divergence.

Existing method, for each iteration, only matches the score function at a single timestep sample  $\hat{z}_t$ , with uniform sampling strategy of time  $t$ . This leads to slow convergence or even difficulty in optimization during Stage 2, as the gradients of important timesteps are averaged out. To this end, we propose our Distribution-Aware Sampling Module (DASM). This module accumulates optimization gradients for earlier timestep samples in a single iteration, enabling the backpropagation of more gradients focused on detail optimization. As shown in Fig. 6, we first introduce noise into the synthetic latent representation  $\hat{z}_t = (1 - \sigma_t)\hat{z}_0 + \sigma_t \epsilon$ , where  $\sigma$  is a weight factor and  $\epsilon$  is a gaussian noise. Subsequently, we employ a LoRA Model to perform denoising, yielding the previous time noisy samples by Eq. (6):

$$\hat{z}_{t-1} = \hat{z}_t + (\sigma_{t-1} - \sigma_t) \cdot \epsilon_\phi(\hat{z}_t; t, c_y), \quad (6)$$

where  $\sigma_{t-1}$  and  $\sigma_t$  are from the flow matching scheduler, and here LoRA Model has fit  $\hat{z}_0$ 's distribution. Similarly,  $\hat{z}_{t-1}$  can be obtained by denoising on the Teacher Model. Ultimately, in a single iteration, the gradients from noise samples along its trajectory can be accumulated to update the Student Model. Since the obtained samples follow the diffusion sampling trajectory and are directed towards early

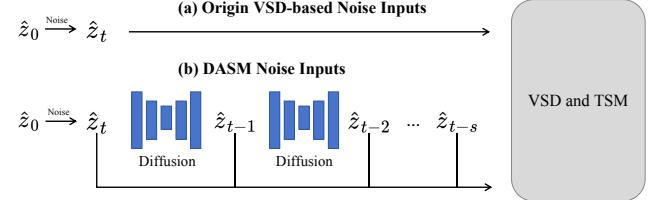


Figure 6. Illustration of DASM. Top: The naive method that adds noise to samples. Bottom: Our DASM leverages the priors of diffusion models to sample noise that more closely aligns with the true sampling trajectory. The noisy samples can all serve as inputs to the subsequent network to generate gradient backpropagation.

timestep, we effectively optimize the divergence in Stage 2.

### 3.5. Training Objective

We summarize all the losses we used in our framework.

**Student Model  $G_\theta$ .** We train our Student Model with the reconstruction loss  $\mathcal{L}_{Rec}$  and the regularization loss  $\mathcal{L}_{Reg}$ . For the reconstruction loss, we use LPIPS loss:

$$\mathcal{L}_{Rec}(G_\theta(\mathbf{x}_L), \mathbf{x}_H) = \mathcal{L}_{LPIPS}(G_\theta(\mathbf{x}_L), \mathbf{x}_H). \quad (7)$$

For the regularization loss, we use our TSD loss Eq. (5). Therefore, the overall training objective for the Student Model  $G_\theta$  is:

$$\mathcal{L}_{Stu} = \mathcal{L}_{Rec} + \gamma \mathcal{L}_{Reg}, \quad (8)$$

where  $\gamma$  is a weighting scalar.

**LoRA Model  $\epsilon_\phi$ .** As stipulated by VSD, the fine-tuned replica  $\epsilon_\phi$  must be trainable, with its training objective being:

$$\mathcal{L}_{Diff}(\hat{z}, c_y) = \mathbb{E}_{t, \epsilon} [\|\epsilon_\phi(\hat{z}_t; t, c_y) - \epsilon'\|^2], \quad (9)$$

where  $\epsilon'$  serves as the training target for the denoising network, representing Gaussian noise in the context of DDPM, and a gradient towards HQ for flow matching.

## 4. Experiments

### 4.1. Experimental Settings

**Training Datasets.** We adopt DIV2K [1], Flickr2K [42], LSDIR [27] and first 10K face images from FFHQ [20] for training. We use the same degradation pipeline as Real-ESRGAN [46] to synthesize LR-HR pairs.

**Test Datasets.** We evaluated our model on synthetic DIV2K-Val [1] dataset and two real-world datasets, RealSR [4] and DRealSR [51]. The real-world datasets comprised 128x128 LQ and 512x512 HQ image pairs.

**Evaluation Metrics.** For evaluating our method, we apply both full-reference and no-reference metrics. Full-reference metrics include PSNR and SSIM [49] (calculated on the Y channel in YCbCr space) for fidelity, LPIPS [65] and DISTS [7] for perceptual quality, and FID [14] for distribution comparison. No-reference metrics include NIQE [63], MANIQA[55], MUSIQ [22] and CLIPQA [43].

Table 1. Quantitative comparison with the state-of-the-art one-step methods across both synthetic and real-world benchmarks. The number of diffusion inference steps is indicated by ‘s’. The best results of each metric are highlighted in red.

Datasets	Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	DISTS $\downarrow$	FID $\downarrow$	NIQE $\downarrow$	MUSIQ $\uparrow$	MANIQA $\uparrow$	CLIPQA $\uparrow$
DRealSR	OSEDiff-1s	27.92	<b>0.7836</b>	0.2968	0.2162	135.51	6.4471	64.69	<b>0.5898</b>	0.6958
	AddSR-1s	27.77	0.7722	0.3196	0.2242	150.18	6.9321	60.85	0.5490	0.6188
	SinSR-1s	<b>28.38</b>	0.7497	0.3669	0.2484	172.72	6.9606	55.03	0.4904	0.6412
	Ours-1s	27.77	0.7559	<b>0.2967</b>	<b>0.2136</b>	<b>134.98</b>	<b>5.9131</b>	<b>66.62</b>	0.5874	<b>0.7344</b>
RealSR	OSEDiff-1s	25.15	0.7341	0.2920	0.2128	123.48	5.6471	69.10	0.6326	0.6687
	AddSR-1s	24.79	0.7077	0.3091	0.2191	132.05	5.5440	66.18	0.6098	0.5722
	SinSR-1s	<b>26.27</b>	<b>0.7351</b>	0.3217	0.2341	137.59	6.2964	60.76	0.5418	0.6163
	Ours-1s	24.81	0.7172	<b>0.2743</b>	<b>0.2104</b>	<b>114.45</b>	<b>5.1298</b>	<b>71.19</b>	<b>0.6347</b>	<b>0.7160</b>
DIV2K-Val	OSEDiff-1s	23.72	<b>0.6109</b>	0.2942	0.1975	<b>26.34</b>	4.7089	67.96	0.6131	0.6681
	AddSR-1s	23.26	0.5902	0.3623	0.2123	29.68	4.7610	63.39	0.5657	0.5734
	SinSR-1s	<b>24.41</b>	0.6018	0.3262	0.2068	35.55	5.9981	62.95	0.5430	0.6501
	Ours-1s	23.02	0.5808	<b>0.2673</b>	<b>0.1821</b>	29.16	<b>4.3244</b>	<b>71.69</b>	<b>0.6192</b>	<b>0.7416</b>

**Compared Methods.** We categorize the test models into two groups: single-step and multi-step inference. The single-step inference diffusion models include SinSR [48], AddSR [54], and OSEDiff [52]. The multi-step inference diffusion models comprise StableSR [44], ResShift [61], PASD [56], DiffBIR [29], SeeSR [53], SUPIR [60], and AddSR [54]. Specifically, for AddSR, we have conducted comparisons between its single-step and four-step models. GAN-based Real-ISR methods [5, 28, 46, 62] are detailed in the supplementary material.

**Implementation Details.** All models mentioned are initialized from the Teacher Model (SD3 [10] in our work). We only train vae encoder and denoising network in the Student Model, frozen the vae decoder in order to preserve vae’s prior [24]. We utilize the default prompt for the Student Model, and prompts extracted from HQ for the Teacher and LoRA models when training. We use the AdamW optimizer [30] with a learning rate of 5e-6 for the Student Model and 1e-6 for the LoRA Model, and set rank of LoRA to 64 for both two models. The training process took roughly 120h, utilizing 8 NVIDIA V100 GPUs with a batch size of 16.

## 4.2. Comparison with Existing Methods

**Quantitative Comparisons.** Tab. 1 show the quantitative comparison of our method with single-step diffusion models on three datasets. Ours achieves the best results on most evaluation metrics. SinSR and AddSR, as distilled versions of previous multi-step super-resolution methods, reduce inference steps but experience a corresponding decrease in performance metrics. OSEDiff introduces the VSD loss from 3D generation tasks into Real-ISR without fully addressing the substantial differences between the two domains. Therefore, the no-reference metrics for image quality are not very satisfied. In contrast, our proposed TSD-SR, tailored for Real-ISR, outperform other single-step models in terms of the vast majority of key metrics.

Tab. 2 show the quantitative comparison with multi-step models. We can draw the following conclusions: (1) TSD-SR demonstrates significant advantages over competing methods in terms of LPIPS, DISTS, and NIQE metrics. Additionally, it achieves performance that surpasses most multi-step models on FID, MUSIQ, and CLIPQA. (2) DiffBIR, SeeSR, PASD, and AddSR exhibit better performance on the MANIQA metric, which may be attributed to the fact that multi-step models have more denoising iterations to produce rich details. (3) ResShift stands out with the highest PSNR and SSIM scores, while StableSR also shows notable performance regarding DISTS and FID metrics. However, both two models underperform in terms of no-reference metrics.

Finally, we explain the lower PSNR and SSIM metrics in our experiments. Several works [54, 60] have found that these reconstruction metrics are not well-suited for evaluating Real-ISR tasks. When a model recovers better details, it leads to lower reconstruction metrics, indicating a trade-off [3, 32, 67]. Refer to the supplementary material for detailed visual comparisons and analysis.

**Qualitative Comparisons.** Fig. 7 present visual comparisons of different Real-ISR methods. As shown in multi-step methods’ results, SeeSR uses degradation-aware semantic cues to leverage image generation priors, but it sometimes produces over-smooth textures. SUPIR demonstrates a notably robust generative capacity, however, the overproduction of extraneous details will lead to results that lack naturalism in image restoration (e.g. Adding excessive wrinkles at the corners of a young girl’s eyes). Under more realistic degradation conditions, PASD finds it challenging to restore the corresponding content, revealing a deficiency in the model’s robustness capabilities. Among single-step methods, SinSR introduces artifacts, possibly because its distilled pre-trained model is trained from scratch, lacking sufficient real-world priors, which results in unsatisfac-

Table 2. Quantitative comparison with state-of-the-art multi-step methods across both synthetic and real-world benchmarks. The number of diffusion inference steps is indicated by ‘s’. The best and second best results of each metric are highlighted in red and blue, respectively.

Datasets	Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	DISTS $\downarrow$	FID $\downarrow$	NIQE $\downarrow$	MUSIQ $\uparrow$	MANIQA $\uparrow$	CLIPQA $\uparrow$
DRealSR	StableSR-200s	28.04	0.7454	0.3279	<b>0.2272</b>	<b>144.15</b>	6.5999	58.53	0.5603	0.6250
	DiffBIR-50s	25.93	0.6525	0.4518	0.2761	177.04	<b>6.2324</b>	65.66	<b>0.6296</b>	0.6860
	SeeSR-50s	<b>28.14</b>	<b>0.7712</b>	<b>0.3141</b>	0.2297	146.95	6.4632	64.74	0.6022	0.6893
	SUPIR-50s	25.09	0.6460	0.4243	0.2795	169.48	7.3918	58.79	0.5471	0.6749
	PASD-20s	27.79	0.7495	0.3579	0.2524	171.03	6.7661	63.23	0.5919	0.6242
	ResShift-15s	<b>28.69</b>	<b>0.7874</b>	0.3525	0.2541	176.77	7.8762	52.40	0.4756	0.5413
	AddSR-4s	26.72	0.7124	0.3982	0.2711	164.12	7.6689	<b>66.33</b>	<b>0.6257</b>	<b>0.7226</b>
	Ours-1s	27.77	0.7559	<b>0.2967</b>	<b>0.2136</b>	<b>134.98</b>	<b>5.9131</b>	<b>66.62</b>	0.5874	<b>0.7344</b>
RealSR	StableSR-200s	24.62	0.7041	0.3070	<b>0.2156</b>	128.54	5.7817	65.48	0.6223	0.6198
	DiffBIR-50s	24.24	0.6650	0.3469	0.2300	134.56	5.4932	68.35	<b>0.6544</b>	0.6961
	SeeSR-50s	25.21	0.7216	<b>0.3003</b>	0.2218	<b>125.10</b>	<b>5.3978</b>	69.69	0.6443	0.6671
	SUPIR-50s	23.65	0.6620	0.3541	0.2488	130.38	6.1099	62.09	0.5780	0.6707
	PASD-20s	<b>25.68</b>	<b>0.7273</b>	0.3144	0.2304	134.18	5.7616	68.33	0.6323	0.5783
	ResShift-15s	<b>26.39</b>	<b>0.7567</b>	0.3158	0.2432	149.59	6.8746	60.22	0.5419	0.5496
	AddSR-4s	23.33	0.6400	0.3925	0.2626	154.22	5.8959	<b>71.49</b>	<b>0.6826</b>	<b>0.7225</b>
	Ours-1s	24.81	0.7172	<b>0.2743</b>	<b>0.2104</b>	<b>114.45</b>	<b>5.1298</b>	<b>71.19</b>	0.6347	<b>0.7160</b>
DIV2K-Val	StableSR-200s	23.27	0.5722	<b>0.3111</b>	0.2046	<b>24.95</b>	4.7737	65.78	0.6164	0.6753
	DiffBIR-50s	23.13	0.5717	0.3469	0.2108	33.93	<b>4.6056</b>	68.54	<b>0.6360</b>	0.7125
	SeeSR-50s	23.73	<b>0.6057</b>	0.3198	<b>0.1953</b>	<b>25.81</b>	4.8322	68.49	0.6198	0.6899
	SUPIR-50s	22.13	0.5279	0.3919	0.2312	31.40	5.6767	63.86	0.5903	0.7146
	PASD-20s	<b>24.00</b>	0.6041	0.3779	0.2305	39.12	4.8587	67.36	0.6121	0.6327
	ResShift-15s	<b>24.71</b>	<b>0.6234</b>	0.3473	0.2253	42.01	6.3615	60.63	0.5283	0.5962
	AddSR-4s	22.16	0.5280	0.4053	0.2360	35.41	5.2584	<b>70.99</b>	<b>0.6596</b>	<b>0.7593</b>
	Ours-1s	23.02	0.5808	<b>0.2673</b>	<b>0.1821</b>	29.16	<b>4.3244</b>	<b>71.69</b>	0.6192	<b>0.7416</b>

Table 3. Comparison of computational complexity across different DMs-based methods. The performance of each method is measured on an A100 GPU, using input images sized at  $512 \times 512$  pixels. We have disregarded the loading time for model weights and data.

	StableSR	DiffBIR	SeeSR	SUPIR	PASD	ResShift	AddSR	AddSR	SinSR	OSEDiff	Ours
Inference Step	200	50	50	50	20	15	4	1	1	1	1
Inference Time	12.4151	7.9637	5.8167	16.8704	4.8441	0.7546	1.0199	0.5043	<b>0.1424</b>	0.1650	<b>0.1362</b>

tory image restoration capabilities. AddSR produces over-smoothed results when using its 1-step model. OSEDiff offers improved restoration effects compared to SinSR and AddSR, yet it may fall short in terms of authenticity and naturalness when it comes to detail recovery. In contrast, our method effectively generates rich textures and realistic details with enhanced sharpness and contrast. Additional visual comparisons and results are provided in the **supplementary material**.

**Complexity Comparisons** We assess the complexity of the state-of-the-art DM-based Real-ISR models as detailed in Tab. 3, focusing on inference time. The performance of each method is benchmarked on an A100 GPU with input images sized at  $512 \times 512$  pixels. We have disregarded the loading time for model weights and data, and the main computation time includes: (1) text extraction time when a text extractor is used; (2) text encoder computation time if needed; (3) VAE encoding and decoding time; (4) denoising network execution time. It is evident that TSD-SR has a substantial advantage in inference time compared with

multi-step models. Specifically, TSD-SR is more than 120 times faster than SUPIR, 90 times faster than StableSR, about 50 times faster than DiffBIR, more than 40 times faster than SeeSR, more than 35 times faster than PASD, and 4 times faster than ResShift. When compared with existing one-step models, our model boasts the fastest inference times. This is because we directly denoise from LQ data and use a fixed prompt.

### 4.3. User Study

We conduct a user study comparing our method with three other diffusion-based one-step super-resolution methods. To make the evaluation more comprehensive, we selected images from five categories—human faces, buildings, animals, vegetation and characters. A total of 50 participants were engaged in the voting process. We guided participants to evaluate the best restoration results based on the similarity to HQ image, structural similarity to the LQ image, and realism of textures and details. The results shown in the Fig. 8 indicate that our method receives a 69.2% approval

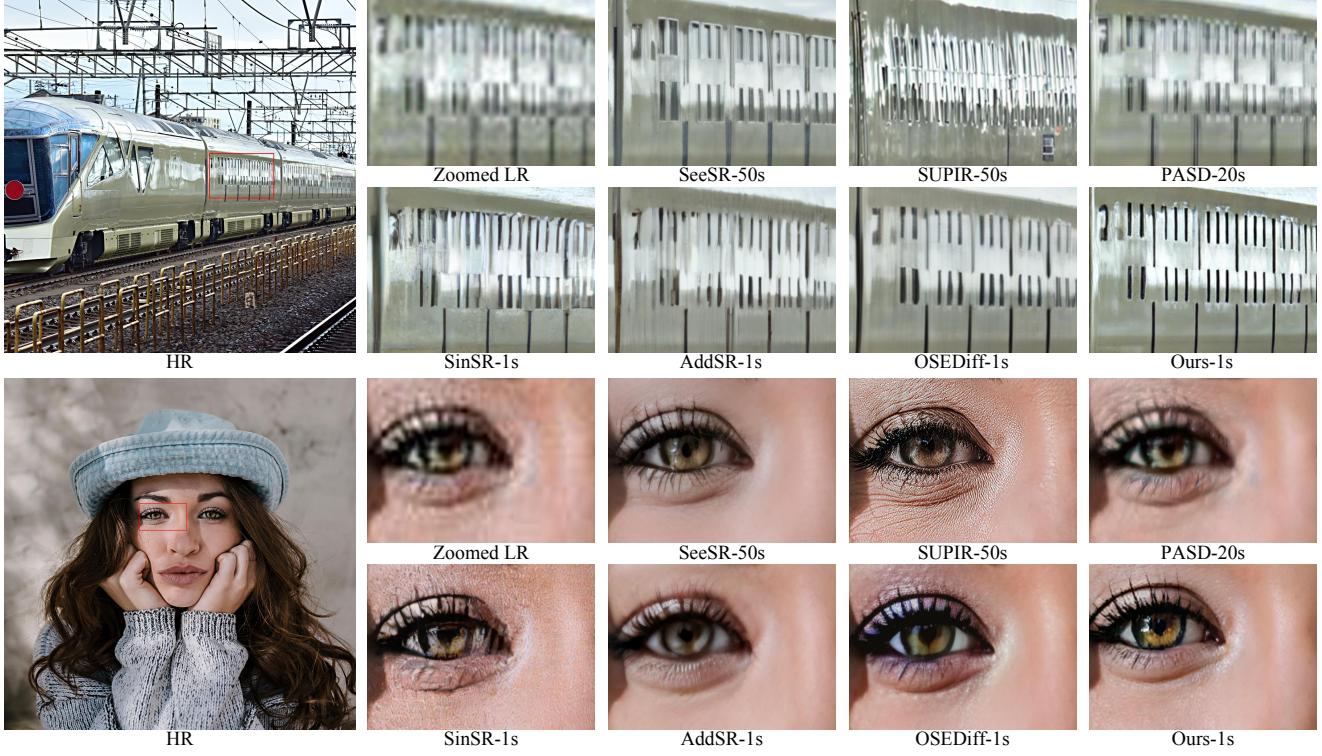


Figure 7. Visual comparisons of different Real-ISR methods. Please zoom in for a better view.



Figure 8. The results of our user study. Left: Category-based user preference radar chart. Our model has won the most user favor across all categories. Right: User preference pie chart. Our approach has garnered a 69.2% user satisfaction rating.

Table 4. Ablation study of Target Score Matching loss and Distribution-Aware Sampling Module.

Datasets	Method	DISTS $\downarrow$	MUSIQ $\uparrow$	MANIQA $\uparrow$	CLIPQA $\uparrow$
DRealSR	w/o TSM	0.2327	63.90	0.5749	0.6958
	w/o DASM	0.2311	63.56	0.5812	0.7123
	Full	<b>0.2136</b>	<b>66.62</b>	<b>0.5874</b>	<b>0.7344</b>
RealSR	w/o TSM	0.2397	68.56	0.6338	0.6987
	w/o DASM	0.2273	69.05	0.6273	0.7031
	Full	<b>0.2104</b>	<b>71.19</b>	<b>0.6347</b>	<b>0.7160</b>

rate from users. We scored 57.60% in Animals, 70.00% in Buildings, 68.80% in Human Faces, 65.20% in Vegetation, and 84.40% in Characters, surpassing other models.

#### 4.4. Ablation Study

To validate the effectiveness of TSM loss and DASM, we conduct ablation studies by removing them separately in

our experiments. We select DISTS, MUSIQ, MANIQA and CLIPQA metrics for comparison, as these are critical for image quality assessment. The results are presented in Tab. 4. We draw the following conclusions: (1) The absence of TSM loss and DASM affects the metrics DISTS, MUSIQ, MANIQA and CLIPQA. (2) The lack of TSM leads to a significant decrease in DISTS, MUSIQ and CLIPQA metrics, possibly due to the unreliable directions in VSD causing unrealistic generations. (3) The absence of DASM results in a decline in MUSIQ and CLIPQA, possibly due to the suboptimal optimization of details.

## 5. Conclusion and Limitation

We propose TSD-SR, an effective and one-step model based on diffusion prior for Real-ISR. TSD-SR utilizes the TSD to enhance the realism of images generated by the distillation model. And it leverages DASM to sample distribution-based samples and accumulate their gradients to enhance the recovery of details. Our experiments have demonstrated that TSD-SR outperforms existing one-step Real-ISR models in both performance and inference speed.

**Limitations.** Although our model boasts excellent inference speed and restoration performance, it still has a large number of model parameters compared to the past GAN/non-diffusion approaches. In the future, we plan to employ pruning or quantization methods to compress the model parameters, striving for a lightweight and efficient Real-ISR model.

**Acknowledgment.** This work was supported by National Natural Science Foundation of China (62293554, U2336212), Zhejiang Provincial Natural Science Foundation of China (LD24F020007), National Key R&D Program of China (SQ2023AAA01005), Ningbo Innovation “Yongjiang 2035” Key Research and Development Programme (2024Z292), and Young Elite Scientists Sponsorship Program by CAST (2023QNRC001).

## References

[1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 126–135, 2017. 5

[2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017. 1, 2

[3] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6228–6237, 2018. 6

[4] 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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3086–3095, 2019. 5

[5] Chaofeng Chen, Xinyu Shi, Yipeng Qin, Xiaoming Li, Xiaoguang Han, Tao Yang, and Shihui Guo. Real-world blind super-resolution via feature matching with implicit high-resolution priors. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1329–1338, 2022. 6

[6] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12299–12310, 2021. 1

[7] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence*, 44(5):2567–2581, 2020. 2, 5

[8] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*, pages 184–199. Springer, 2014. 1

[9] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015. 1

[10] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 2, 6

[11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 1

[12] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021. 2

[13] Xiangyu He and Jian Cheng. Revisiting l1 loss in super-resolution: a probabilistic view and beyond. *arXiv preprint arXiv:2201.10084*, 2022. 4

[14] 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. *Advances in neural information processing systems*, 30, 2017. 5

[15] Geoffrey Hinton. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2

[16] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 2

[17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1

[18] Andrew G Howard. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 2

[19] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 3

[20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 5

[21] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022. 1

[22] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5148–5157, 2021. 5

[23] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016. 1

[24] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 6

[25] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951. 3

[26] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on*

*computer vision and pattern recognition*, pages 4681–4690, 2017. 1, 2

[27] Yawei Li, Kai Zhang, Jingyun Liang, Jiezhang Cao, Ce Liu, Rui Gong, Yulun Zhang, Hao Tang, Yun Liu, Denis Deman-dolx, et al. Lsdir: A large scale dataset for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1775–1787, 2023. 5

[28] Jie Liang, Hui Zeng, and Lei Zhang. Details or artifacts: A locally discriminative learning approach to realistic image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5657–5666, 2022. 6

[29] 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. *arXiv preprint arXiv:2308.15070*, 2023. 1, 2, 6

[30] I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6

[31] Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*, 2023. 2

[32] Xiaotong Luo, Yuan Xie, Yanyun Qu, and Yun Fu. Skipdiff: Adaptive skip diffusion model for high-fidelity perceptual image super-resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4017–4025, 2024. 6

[33] Mehdi Mirza. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 1

[34] Thuan Hoang Nguyen and Anh Tran. Swiftbrush: One-step text-to-image diffusion model with variational score distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7807–7816, 2024. 2

[35] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 2

[36] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 3

[37] Alec Radford. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 1

[38] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1 (2):3, 2022. 2

[39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 2

[40] Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation. *arXiv preprint arXiv:2311.17042*, 2023. 2

[41] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 1

[42] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, and Lei Zhang. Ntire 2017 challenge on single image super-resolution: Methods and results. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 114–125, 2017. 5

[43] 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*, pages 2555–2563, 2023. 5

[44] Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin CK Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. *International Journal of Computer Vision*, pages 1–21, 2024. 2, 6

[45] 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 *Proceedings of the European conference on computer vision (ECCV) workshops*, pages 0–0, 2018. 2, 3

[46] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1905–1914, 2021. 1, 2, 3, 5, 6

[47] Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. *arXiv preprint arXiv:2212.00490*, 2022. 1

[48] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25796–25805, 2024. 2, 6

[49] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 5

[50] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in Neural Information Processing Systems*, 36, 2024. 2, 3

[51] Pengxu Wei, Ziwei Xie, Hannan Lu, Zongyuan Zhan, Qixiang Ye, Wangmeng Zuo, and Liang Lin. Component divide-and-conquer for real-world image super-resolution. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII 16*, pages 101–117. Springer, 2020. 5

[52] Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. *arXiv preprint arXiv:2406.08177*, 2024. 2, 3, 4, 6

[53] Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. Seesr: Towards semantics-aware real-world image super-resolution. In *Proceedings of*

*the IEEE/CVF conference on computer vision and pattern recognition*, pages 25456–25467, 2024. 1, 2, 6

[54] Rui Xie, Ying Tai, Kai Zhang, Zhenyu Zhang, Jun Zhou, and Jian Yang. Addsr: Accelerating diffusion-based blind super-resolution with adversarial diffusion distillation. *arXiv preprint arXiv:2404.01717*, 2024. 2, 6

[55] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1191–1200, 2022. 5

[56] Tao Yang, Rongyuan Wu, Peiran Ren, Xuansong Xie, and Lei Zhang. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. *arXiv preprint arXiv:2308.14469*, 2023. 1, 2, 6

[57] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4133–4141, 2017. 2

[58] Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and William T Freeman. Improved distribution matching distillation for fast image synthesis. *arXiv preprint arXiv:2405.14867*, 2024. 2

[59] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6613–6623, 2024. 2

[60] Fanghua Yu, Jinjin Gu, Zheyuan Li, Jinfan Hu, Xiangtao Kong, Xintao Wang, Jingwen He, Yu Qiao, and Chao Dong. Scaling up to excellence: Practicing model scaling for photo-realistic image restoration in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25669–25680, 2024. 1, 2, 6

[61] Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Resshift: Efficient diffusion model for image super-resolution by residual shifting. *Advances in Neural Information Processing Systems*, 36, 2024. 2, 6

[62] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4791–4800, 2021. 1, 2, 3, 6

[63] Lin Zhang, Lei Zhang, and Alan C Bovik. A feature-enriched completely blind image quality evaluator. *IEEE Transactions on Image Processing*, 24(8):2579–2591, 2015. 5

[64] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 2

[65] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 2, 3, 5

[66] Xindong Zhang, Hui Zeng, Shi Guo, and Lei Zhang. Efficient long-range attention network for image super-resolution. In *European conference on computer vision*, pages 649–667. Springer, 2022. 1

[67] Yuehan Zhang, Bo Ji, Jia Hao, and Angela Yao. Perception-distortion balanced admm optimization for single-image super-resolution. In *European Conference on Computer Vision*, pages 108–125. Springer, 2022. 6