AdaDiffSR: Adaptive Region-aware Dynamic Acceleration Diffusion Model for Real-World Image Super-Resolution

Yuanting Fan^{1,2*}^o, Chengxu Liu^{2,3*}^o, Nengzhong Yin², Changlong Gao², and Xueming Qian^{2,3⊠}^o

¹ School of Software Engineering, Xi'an Jiaotong University, Xi'an, China
² Xi'an Jiaotong University, Xi'an, China

 $^3\,$ Shaanxi Yulan Jiuzhou Intelligent Optoelectronic Tech. Co., Ltd, Xi'an, China

{retofan,chengxuliu,ynz0608,gaochanglong}@stu.xjtu.edu.cn,

qianxm@mail.xjtu.edu.cn

Abstract. Diffusion models (DMs) have shown promising results on single-image super-resolution and other image-to-image translation tasks. Benefiting from more computational resources and longer inference times, they are able to yield more realistic images. Existing DMs-based superresolution methods try to achieve an overall average recovery over all regions via iterative refinement, ignoring the consideration that different input image regions require different timesteps to reconstruct. In this work, we notice that previous DMs-based super-resolution methods suffer from wasting computational resources to reconstruct invisible details. To further improve the utilization of computational resources, we propose AdaDiffSR, a DMs-based SR pipeline with dynamic timesteps sampling strategy (DTSS). Specifically, by introducing the multi-metrics latent entropy module (MMLE), we can achieve dynamic perception of the latent spatial information gain during the denoising process, thereby guiding the dynamic selection of the timesteps. In addition, we adopt a progressive feature injection module (PFJ), which dynamically injects the original image features into the denoising process based on the current information gain, so as to generate images with both fidelity and realism. Experiments show that our AdaDiffSR achieves comparable performance over current state-of-the-art DMs-based SR methods while consuming less computational resources and inference time on both synthetic and real-world datasets.

Keywords: Super resolution · Diffusion models · Adaptive inference

1 Introduction

Real-world image super-resolution aims to recover a realistic high-resolution (HR) image from an unknown degraded low-resolution (LR) counterpart. It is adopted to enhance the image visual quality and is widely used in the fields of

^{*} Equal contribution.



Fig. 1: Visual comparisons between background and foreground regions during denoising process. The red and blue boxes represent the background and foreground regions, respectively. We visualize the variations of several corresponding metrics as the radar chart on the right, the smaller the better. As the timesteps increase from 50 to 200, we find that the visual results of the foreground regions become more satisfactory while the background remains almost unchanged.

satellite imagery [26], surveillance, etc. Nowadays, significant advances have been made in diffusion models for image synthesis tasks [14,18,40,59]. Recent research demonstrates that the pre-trained models (*e.g.*, SD [41]) can easily transferred to various downstream low-level vision tasks, including image and video translation [41–43, 51]. In this work, we explore the potential of using pre-trained diffusion models to perform super-resolution efficiently and effectively.

According to whether using pre-trained parameters, Existing DMs-based SR methods [41-43,51] can be divided into two paradigms. The former [42,43] trains the DMs-based SR model from scratch while using LR features as additional input to constrain the fidelity of output. The latter [41, 51] incorporates constraints into the denoising progress of the pre-trained diffusion models, which not only avoids the demand for computational resources to train the diffusion models, but also enables better utilization of the diffusion prior within the pretrained diffusion models. However, all of the above methods inherently fail to take into account the relationship between reconstructing different image regions and the computational resources they cost. It means that some regions may require fewer timesteps to achieve a satisfactory reconstruction performance, while others require larger timesteps. Consequently, these methods exhibit sub-optimal utilization of computational resources. As shown in Fig. 1, with the increasing of timesteps in the denoising process, the background regions (e.q., sky) or sea in red boxes) are not visually distinct, and only the foreground regions (e.g.,surfaces or structures in blue boxes) become more realistic.

In this paper, we propose a novel DMs-based SR pipeline, dubbed AdaDiffSR. It utilizes the latent information gain to dynamically adjust timesteps during inference and effectively reduces the computational resource overhead. Specifically, we divide the entire input image into a series of sub-regions separately for the denoising process. For these sub-regions, as shown in Fig. 2, we design a dynamic timesteps sampling strategy (DTSS) that adjusts timesteps dynamically according to the information gain within these sub-regions. To implement it, we propose a multi-metrics latent entropy (MMLE) module, which incorporates multiple image quality assessment (IQA) metrics and can perceive information gain dynamically from both qualitative and quantitative perspectives. In addition, to balance the fidelity of the reconstructed image with the texture generation capability of diffusion models, we further propose a progressive feature injection (PFJ) module. It dynamically adjusts the fusion process between original input images and denoising features according to the current information gain. Meanwhile, by applying the latent space region integration strategy, discontinuities between different restored sub-regions are effectively eliminated.

Our main contributions are as follows:

- We introduce AdaDiffSR, a diffusion model for real-world image superresolution. It is one of the first works to focus on optimizing the utilization of computational resources in DMs-based SR paradigms and provides a new inspiration for future research on denoising models.
- We design a dynamic timesteps sampling strategy (DTSS), in conjunction with the proposed multi-metrics latent entropy (MMLE) module, which can significantly reduce the computational overhead of the denoising process while obtaining superior results.
- We propose a progressive feature injection (PFJ) module to dynamically incorporate original image features into the denoising process, allowing the recovered results with fidelity and rich texture details.
- Extensive experiments demonstrate that the proposed AdaDiffSR can significantly reduce computational resources and inference time compared to existing SOTA methods.

2 Related Work

2.1 Single Image Super-Resolution

Single Image Super-Resolution (SISR) aims to restore the corresponding HR counterpart from the degraded LR image. The previous SR methods [10, 11, 24, 26, 27] focused on learning the mappings between LR and HR image for particular degradation kernel (*e.g.*, Bicubic downsampling or Gaussian blur kernel). However, the above mappings have limited generalization capacity in real-world scenarios and cannot achieve better visual performance.

Recently, more works [8, 12, 29–31, 34, 53, 60, 70] have focused on optimizing the performance of real-world SR (or blind SR), which aims to learn the similar degradation process in real-world scenarios. Due to the lack of paired real-world training datasets, some methods [55, 71] explicitly synthesize the paired LR-HR images for blind SR. Specifically, recent works [55, 71] present the effective random degradation process, which produces the LR counterparts by applying blur before downsampling the HR images, and then adding noise and applying JPEG compression to the downsampled results. These degradation schemes have 4 Y. Fan et al.

been crucial for GAN-based SR methods to achieve state-of-the-art performance. Building upon these schemes, recent DMs-based SR methods [42, 43, 51, 67] further show more satisfactory performance in real-world scenarios. However, the above SR methods failed to consider the utilization of computational resources during the reconstruction process. In this work, we consider the latent information gain and adjust timesteps of different image regions dynamically during the denoising process, thus reducing the computational resource overhead effectively.

2.2 Diffusion Model in Super-Resolution

With the rapid development of diffusion models (DMs) in the image and video generation tasks [59], recently, numerous DMs-based SR methods [28, 36, 42, 43, 51, 67] have been proposed. Compared to the previous GAN-based SR methods [55, 57, 70] suffer from mode collapse [4, 35, 45] and convergence difficulty during the training process, the existing DMs-based SR methods obtain the SR results via iteratively refine the original noise images, which significantly enhances the model robustness and generalization capacity. Specifically, SR3 [43] first introduces the diffusion models into the SR tasks and achieves state-of-the-art performance on the face and natural datasets. Recent works [42, 51] have attempted to improve DMs-based SR methods from the perspectives of image degradation and image fidelity. In this work, we consider the utilization of computational resources in DMs-based SR methods and propose to dynamically adjust the timesteps of different image regions during the denoising process.

2.3 Adaptive Inference

The previous methods [64–66] focus more on designing slimmable network structures, which adaptively approximate the network performance according to the data characteristics. Specifically, AutoSlim [64] trains a single slimmable network to adjust the network performance of different channel configurations, and then searches the optimized channel configurations under different resources.

Instead of adjusting the network structure, inspired by the fact that different regions have different restoration difficulties, ClassSR [22] utilizes the characteristics of different image regions by dividing the images into patches, then apply different capacity models to the corresponding restoration difficulties. However, the application of multiple models leads to a significant increase in model parameters. APE [54] proposes to learn the incremental capacity of each model layer instead of restoration difficulties of patches, enabling the patch to exit at the optimal layer, thus achieving more practical speedup. Although existing methods [32, 49, 52] have taken into account the information variation (network layer-wise) of different image regions during the reconstruction process, they lack further exploration of the latent information variation (timestep-wise) during the denoising process in diffusion models. In this work, we combine multiple IQA metrics to measure the latent information gain during the denoising process in DMs-based SR methods, thus guiding the dynamic timesteps selection and reducing the computational resource overhead effectively.



Fig. 2: The framework of AdaDiffSR. We calculate the information gain during the denoising process, via information gain, we can modulate the original image features to guide the PFJ module and adjust timesteps dynamically for different regions, thus achieving a trade-off between the computational resources and restoration quality.

3 Methodology

3.1 Motivation

To reduce the computational resource overhead of DMs-based SR methods, we observe that although latent diffusion models can significantly reduce computational resource overhead by reducing feature size, as shown in Fig. 1, there are still a lot of computational resources used to reconstruct invisible details, which leads to sub-optimal utilization of computational resources. Therefore, we propose AdaDiffSR to further improve the utilization of computational resources. As shown in Fig. 2, AdaDiffSR consists of the following three components: multimetrics latent entropy module (MMLE) as shown in Sec. 3.2, dynamic timesteps sampling strategy (DTSS) as shown in Sec. 3.3, and the progressive feature injection module (PFJ) and region integration strategy are as shown in Sec. 3.4.

3.2 Multi-metrics latent entropy module

Based on the previous IQA metrics [2,6,9,13,19,20,23,33,38,39,46,50,61–63,72–74,76], we propose the multi-metrics latent entropy module (MMLE) to measure the information evolution from many different perspectives during denoising process. More precisely, we calculate these metrics in latent space, and we find that many perceptual-oriented IQA metrics fail to reflect the information evolution process accurately. On the contrary, the traditional psnr-oriented metrics accurately perceive the information evolution. The above phenomenon is due to the massive Gaussian noise incorporated during the denoising process.

As shown in Fig. 3(b), we plot the variations of IQA metrics during the denoising process. To make it easier to observe the trends in variations of these metrics, we divide these metrics into two categories (*i.e.*, NR-IQA metrics (No-Reference) and FR-IQA (Full-Reference) metrics) and normalize all metrics between zero and one. Significant differences between different IQA metrics during the denoising process allow us to divide them into three categories of variations. Specifically, the first category of variations positively correlates with the information evolution process, including PSNR, AHIQ, etc. The second category of



Fig. 3: To demonstrate the validity of different IQA metrics during the denoising process. (a), we visualize the feature in perceptual-oriented metric LPIPS and the corresponding denoised image during the restoration process. (b), we plot the variations of several IQA metrics during the denoising process with 200 timesteps using DDPM sampling strategy [15]. Best viewed when disabled some metrics (*e.g.*, TOPIQ [6], CLIP-IQA [50] and etc.), the variations of these metrics are similar with HyperIQA [46].

variations has no significant correlation with the information evolution process, mainly including TOPIQ, CKDN, HyperIQA, etc. The remaining variations are categorized into the third category, which tends to significantly reflect the information evolution in the second half of the denoising process. Since most of the third category of metrics utilize CNN or Transformer [48] to extract features in latent space, and in the first half of the denoising process, the latent features are not clear enough to reflect image quality due to excessive Gaussian noise. As shown in Fig. 3(a), the noise is gradually eliminated, and the latent feature tends to be apparent, allowing these metrics to perceive the information evolution.

Based on the above conclusions, we select several IQA metrics with different perspectives for discriminating the multi-dimensional information gain between two latent features, thus guiding the subsequent dynamic timesteps sampling strategy (DTSS) and progressive feature injection (PFJ) module. Specifically, we define the information gain I_i , which measures the latent representation performance between current timestep i and the previous timestep i - 1:

$$I_i = \sigma(R_i - R_{i-1}),\tag{1}$$

where $\sigma(\cdot)$ represents the tanh function, R_i is the latent representation performance of timestep *i*. Since information gain stems from multi-dimensional metrics, evolution in some dimensions may lead to degradation in other dimensions, through tanh function, we limit the range of I_i to [-1, 1] to perceive this degradation. As shown in Fig. 3(b), FR-IQA metrics are more robust to the information evolution, which is mainly reflected in the gradient continuity of the corresponding curve. On the contrary, NR-IQA metrics suffer from the limitation of without reference image (*i.e.*, original image feature) and cannot perceive the image quality improvement until the second half of the denoising process. To achieve the trade-off between capacity and efficiency, we select four FR-IQA metrics (*i.e.*, PSNR, LPIPS, AHIQ, and NLPD) and two NR-IQA metrics (*i.e.*, BRISQUE and MUSIQ) to measure the information gain:

$$R_i = \sum_{c \in C} \omega_c \times M_c(f_i, o), \tag{2}$$

where f_i and o represent the denoising feature in timestep i and the original image feature, respectively. ω_c denotes the coefficients of different IQA metrics, which are used to balance the influence of different metrics. M_c means the normalized value of the corresponding IQA metric. C consists of the six metrics mentioned above. Note that for the robustness of information evolution, we only calculate the two NR-IQA metrics information gain during the second half of the denoising process. During inference, considering the tremendous time cost of calculating IQA metrics on the fly, we propose to train a lightweight regressor that consists of multiple convolutional layers and global average pooling layers to estimate the multi-dimensional information gain. More analysis of the IQA metrics selection and the structure of regressor can be found in the ablation studies and supplementary materials.

3.3 Dynamic timesteps sampling strategy

Inspired by previous works [22, 54], as shown in Fig. 1, we find that different regions require different computational resources to achieve a trade-off between utilization of computational resources and reconstruction performance. Benefiting from the above MMLE regressor, we propose a dynamic timesteps sampling strategy (DTSS), which improves the utilization of computational resources and adaptively accelerates the DMs-based SR methods.

Specifically, to preserve the prior encapsulated within pre-trained diffusion models, we first crop the input images into multiple sub-regions with the fixed spatial resolution as pre-trained diffusion models. Then, we calculate the multidimensional information gain using the MMLE regressor during the denoising process, thereby guiding the dynamic timesteps selection for each sub-region.

The specific strategy for dynamic timestep selection is as follows. Firstly, we define the information gain threshold and max timesteps of diffusion models as τ and T_{max} , respectively. Then, we generate a codebook to store the sampling-related parameters for skipping from one timestep to another. With this codebook, we can implement the skip-step strategy, which means that we perform multiple original DDPM [15] or DDIM [44] timesteps at once. To improve the inference time and stability of the denoising process, we define four intervals for the skip-step strategy to adjust timesteps dynamically for different regions. The conclusion obtained from Sec. 3.2 indicates that the FR-IQA metrics are robust, while the NR-IQA metrics depend on the image region context. Therefore, we divide these sub-regions into three categories from FR and NR perspectives:

- 8 Y. Fan et al.
- Stable regions: The FR metrics increasing steadily while the NR metrics are almost unchanged, indicating that the information within these regions is not enhanced in perceptual perspectives (*e.g.*, the background regions in Fig. 1). In these regions, we apply larger intervals for the better utilization of computational resources and inference time.
- Growing regions: The FR and NR metrics increasing steadily together (e.g., the foreground regions in Fig. 1), the information within these regions is growing along with the denoising process. Then we apply a smaller interval to enhance the reconstruction quality.
- Saturated regions: We also notice the phenomenon that NR metrics significantly declined with more denoising iterations, which is due to the information within these regions being saturated. Therefore we save the best results of these regions in NR perspective and exit the denoising process in advance.

Through the information gain threshold τ , we are able to distinguish between different sub-region categories. Note that the category of sub-regions may change during the denoising process, so the interval of skip-step can be modified dynamically. Finally, when the timesteps reach the value of T_{max} , the reconstruction is finished and the final equivalent timesteps are much smaller than the T_{max} , thus achieving adaptive acceleration while keeping restoration quality. More analysis about interval settings can be found in the ablation studies.

3.4 Trade-off between fidelity and realism

Benefiting from the powerful generative capacity of diffusion models, the reconstructed images tend to be more realistic. Nevertheless, similar to other generative model based SR methods, the results may lack similarity from the original input (*i.e.*, lack of fidelity). Therefore, we propose the progressive feature injection module (PFJ) to achieve the trade-off between fidelity and realism. The PFJ module dynamically incorporates latent features of original input images into the denoising process based on the current information gain, thus enhancing fidelity while maintaining realism. Specifically, the PFJ module dynamically regulates the fusion intensity of the original image features according to the information gain from different perspectives:

$$\hat{o} = \alpha \times o + \beta; \quad \alpha, \beta = \phi(o, I_i),$$
(3)

where α and β represent the modulation coefficients of the original image features o. And we design a small CNN network $\phi(\cdot)$ to estimate the modulation coefficients from the current timestep information gain I_i and original image features o. Generally, if the NR perspectives information gain is more prominent, which means that the realism of the current timestep denoising feature is increasing. To balance the fidelity and realism, the $\phi(\cdot)$ needs to predict larger modulation coefficients to adjust the weights of the original image features for better fidelity, and vice versa. Although the problem of image fidelity has been solved, the fixed input resolution of the pre-trained diffusion model limits the further application of DMsbased SR methods. The previous CNN-based SR methods used the naive strategy to integrate the overlapping regions, leading to discontinuities in the boundaries. Inspired by [18,51], we integrate these regions in latent space using the Gaussian weight maps that are generated from Gaussian kernel, thus avoiding the discontinuities in pixel space, as shown in Fig. 6. To speed up the reconstruction while eliminating the mutual noise influence in adjacent regions, we only employ the above integration strategy in the final timestep. The details of the region integration strategy can be found in the supplementary materials.

4 Experiments

4.1 Implementation Details

Following previous work [51], we finetune the pre-trained diffusion model (Stable Diffusion 2.1-base [41]) using the synthetic training dataset. We follow Stable Diffusion to use Adam optimizer and set the learning rate as 5×10^{-5} . To avoid affecting the prior within the pre-trained diffusion model, we conduct the fine-tuning process on 512×512 resolutions. During inference, we set the max timesteps T_{max} and information gain threshold as 1,000 and 5×10^{-3} , respectively. To handle the images with arbitrary resolutions, we adopt the region integration strategy mentioned in Sec. 3.4. To train the MMLE regressor, we follow the Real-ESRGAN [55] degradation pipeline to generate synthetic LR-HR pairs with 512×512 resolutions. Then, we freeze the parameters of the finetuned diffusion model to train the regressor using the L2 loss function.

4.2 Experimental Settings

Datasets. We follow the degradation pipeline of Real-ESRGAN [55] to synthesize LR-HR paired images on DIV2K [1], Flickr2K [47], and OutdoorSceneTraining datasets [56] for our training datasets. Then, we evaluate our method on both synthetic and real-world testsets. For synthetic data, we follow the above degradation strategy and generate LR-HR pairs from the DIV2K validation set [1]. For real-world datasets, we follow the common settings to conduct comparisons on RealSR [3], DRealSR [58], and DPED-iPhone [16] testsets.

Evaluation Metrics For paired testsets (*i.e.*, DIV2K Valid [1], RealSR [3], and DRealSR [58]), follow previous works, we employ several perceptual-based metrics including LPIPS [73], AHIQ [23], and MUSIQ [20] to measure the image perceptual quality. We also employ PSNR and SSIM metrics to indicate differences between different methods. For datasets without ground-truth image (*i.e.*, DPED-iphone [16]), follow previous works, we employ NR-IQA metrics including BRISQUE [37], TOPIQ [6], NIQE [39], and MUSIQ [20] for perceptual quality.

10 Y. Fan et al.

Table 1: Quantitative comparisons with the state-of-the-art methods on both synthetic and real-world datasets. The best and second-best performances are in red and <u>blue</u> color, respectively. Real-ESRGAN+ [55] is abbreviated as R-ESRG.

Datasets	Metrics	RealSR [17]	DASR [25]	BSRGAN [70]	R-ESRG [55]	FeMaSR [7]	LDM [41]	StableSR [51]	ResShift [67]	Ours
	$\mathrm{PSNR}\uparrow$	25.11	28.33	25.26	24.84	22.97	20.58	23.83	24.53	24.25
DIV9V	SSIM \uparrow	0.7170	0.8091	0.7325	0.7287	0.6857	0.5590	0.7059	0.7323	<u>0.7355</u>
Valid [47]	$\text{LPIPS}\downarrow$	0.2134	0.2892	0.2364	0.2284	0.2177	0.2556	0.2328	0.4406	0.2153
vanu [47]	$AHIQ\uparrow$	0.3746	0.4303	0.4047	0.3927	0.3841	0.3761	<u>0.4117</u>	0.3105	0.3985
	$\rm MUSIQ\uparrow$	64.73	55.57	67.42	64.65	67.33	<u>68.05</u>	66.73	67.84	68.81
	$PSNR\uparrow$	25.56	25.87	24.70	24.33	23.58	22.26	23.55	24.79	24.19
DeelCD	SSIM \uparrow	0.7390	0.7560	0.7427	0.7456	0.7132	0.6462	0.7461	0.7423	0.7485
realsr.	$\rm LPIPS \downarrow$	0.3719	0.3832	0.2713	0.2869	0.3016	0.3288	0.2543	0.2524	0.2595
[၁]	$AHIQ\uparrow$	0.2370	0.2709	0.3167	0.3009	0.2965	0.2925	0.3160	0.3061	0.3145
	$\mathrm{MUSIQ}\uparrow$	41.69	35.19	64.62	57.55	60.53	<u>61.42</u>	59.92	59.67	60.47
	$PSNR\uparrow$	27.79	27.96	26.18	25.82	24.56	23.39	24.85	<u>27.87</u>	25.67
DBaalSD	SSIM \uparrow	0.8158	0.8357	0.7934	0.7987	0.7362	0.6723	0.8326	0.8056	0.8415
Drealsr [Eo]	$\text{LPIPS}\downarrow$	0.3851	0.3933	0.2929	0.2818	0.3374	0.3860	0.2853	0.5408	0.2627
[99]	$AHIQ\uparrow$	0.2606	0.2887	0.3502	0.3478	0.3192	0.2859	0.3489	0.2849	0.3496
	$\mathrm{MUSIQ}\uparrow$	22.41	21.88	35.50	35.25	31.78	37.98	35.39	42.68	<u>38.66</u>
DPED- iPhone [16]	BRISQUE \downarrow	<u>5.69</u>	46.95	15.89	16.49	3.46	22.12	13.31	13.43	11.19
	$\rm MUSIQ\uparrow$	51.27	39.81	51.65	50.99	57.19	56.59	49.97	46.87	51.84
	TOPIQ \uparrow	0.4675	0.3365	0.4879	0.4646	0.5439	0.4329	0.4437	0.4611	0.4925
	$\mathrm{NIQE}\downarrow$	3.20	6.19	3.37	3.17	5.09	5.56	3.80	5.58	3.09

4.3 Comparisons with Existing Methods

To verify the effectiveness of our method, we compare AdaDiffSR with several state-of-the-art methods, including RealSR [17], DASR [25], BSRGAN [70], Real-ESRGAN+ [55], FeMaSR [7], LDM [41], StableSR [51], and ResShift [67]. For fair comparisons, we obtain the performance from the official code and models for testing. Note that the metrics reported for LDM [41] in Tab. 1 are obtained on the crop patches of corresponding datasets, because the LDM only supports static input resolution of 512×512 .

Quantitative Comparisons. We first conduct the quantitative comparisons with several state-of-the-art methods on the synthetic dataset [47] and three real-world benchmarks [3, 16, 58]. As shown in Tab. 1, the proposed AdaDiffSR outperforms state-of-the-art methods on several perceptual-based metrics including LPIPS and NIQE. Specifically, on DRealSR, AdaDiffSR achieves a 0.2627 LPIPS score, which is 6.8% better than Real-ESRGAN+ and 25.5% better than other GAN-based SR approaches. Moreover, AdaDiffSR achieves comparable performance with other state-of-the-art methods in other benchmarks. Although RealSR and DASR achieve higher PSNR and SSIM metrics, they show relatively lower performance on perceptual-based metrics. As shown in Fig. 4, they fail to reconstruct more realistic visual details, thus generating inferior visual results.



Fig. 4: Qualitative comparisons on several real-world images. (The first three rows are arbitrary resolutions, and the resolution of the last three rows is fixed at 512×512)

Qualitative Comparisons. To further compare the visual qualities of different approaches, we present visual results on real-world examples from real-world benchmarks and the internet in Fig. 4. Due to the powerful generative capability of diffusion models, AdaDiffSR outperforms CNN-based and GAN-based methods in detail generation. In addition, it can be noticed that AdaDiffSR has significant improvements in realistic details compared to DMs-based methods, especially for surfaces and structures. For example, in the first and fourth rows in Fig. 4, AdaDiffSR significantly enhances the texture details of building surfaces. More visual results can be found in the supplementary materials.

Complexity Comparisons. As for the complexity of AdaDiffSR, Tab. 2 shows the detailed computational resources and inference time between different DMs-based SR methods. #P and RT indicate parameters and runtime, respectively. For fair comparisons, we use the author-released models to report inference times and FLOPs for LR images with different resolutions running on the NVIDIA Tesla 32G-V100 GPUs. All approaches use the same equivalent timesteps as 50 in DDIM sampling strategy [44] and input resolution as 512×512 or 1024×1024 . Compared to the state-of-the-art SR approach StableSR, AdaDiffSR achieves

12 Y. Fan et al.

Table 2: The complexity compar-isons between DMs-based methods.

Table 3: Ablation studies of interval in skipstep on DRealSR dataset [58].

Methods	#P(M) l	FLOPs(G)	$\operatorname{RT}(s)$	Exp	. interval	PSNR 1	$SSIM \uparrow$	LPIPS \downarrow	RT(s)
LDM (2022) [41]	113.60	888.9	<u>6.1</u> -	(a)	(5,5,5,5)	25.43	0.8437	0.2578	13.4
StableSR (2023) [51]	140.91	$1004.2 \\ 4016.8$	7.2 25.9	(b)	(5,10,10,15) (5,10,15,25)	25.57 25.74	$\frac{0.8419}{0.8261}$	$\frac{0.2607}{0.2754}$	$11.8 \\ 8.9$
ResShift (2024) [67]	118.59	922.3 3689.2	19.7 71.3	(d)	(5,15,20,25) (5,15,20,25)	$\frac{25.14}{25.82}$	0.8193	0.2796	<u>8.4</u>
Ours	120.43	751.4 2768.2	4.8 16.8	(e) Our	(10,15,20,25) s $(5,10,15,20)$	$24.83 \\ 25.67$	$0.7965 \\ 0.8415$	$0.2939 \\ 0.2627$	$\frac{7.3}{9.1}$

Table 4: Ablation studies of IQA metricsused in MMLE on RealSR dataset.

Table 5: Ablation studies of PFJ on DPED-iPhone dataset [16].

Exp.	Com	ponents	PSNR 1	SSIM ↑	LPIPS ↓	MUSIQ ↑	$E_{\rm Exp}$ $\frac{\rm C}{\rm C}$	ompo	nents	MUSIO 1	TOPIO ↑	NIOE
r	\mathbf{FR}	NR			•	· · · ·	C	CCA	PFJ	11001001	10116	тн&ц 4
(a)	\checkmark		25.24	0.7532	0.3628	42.53	(a) v	(32.19	0.3765	7.94
(b)		\checkmark	22.37	0.6872	0.2714	51.23	(b)	\checkmark		42.37	0.4314	4.21
Ours	3 √	\checkmark	<u>24.19</u>	0.7485	0.2595	60.47	Ours		\checkmark	51.84	0.4925	3.09

comparable performance using $1.5 \times$ fewer inference times and $2.7 \times$ fewer FLOPs. We also find that adding the proposed approach to the existing DMs-based methods further improves the inference speed while maintaining reconstruction quality. More relative analysis can be found in the supplementary materials.

4.4 Ablation Study

In this section, we conduct the ablation on the different modules, analyze the visual superiority of progressive feature injection and region integration strategy.

Effect of the interval in skip-step. We first analyze the effectiveness of interval in skip-step. As mentioned in Sec. 3.3, we use different timesteps in different image regions, which is accomplished through skip-step strategy. We conduct several different interval settings to demonstrate the influence of intervals on DRealSR [58] dataset using DDPM sampling strategy. As shown in Tab. 3, increasing the interval leads to a significant improvement in inference time. However, it also brings the degradation in overall visual performance and vice versa. Finally, we set the interval of skip-step as 5, 10, 15, and 20 to achieve a trade-off between utilization of computational resources and restoration quality.

Effect of the metrics in MMLE. As mentioned in Sec. 3.2, we choose four FR-IQA metrics and two NR-IQA metrics to measure the information gain during the denoising process. Here, we further verify the impact of the selected



Fig. 5: Visual comparisons of different configurations in MMLE metrics.



Fig. 6: Visualization results of using region integration strategy.

IQA metrics on RealSR [3] dataset. As shown in Tab. 4, only FR-IQA metrics (*i.e.*, Exp. (a)) or NR-IQA metrics (*i.e.*, Exp. (b)) to guide the denoising process will lead to sub-optimal results. Specifically, only with FR metrics (*e.g.*, PSNR or SSIM) leads to the denoising process focusing on the restoration of latent features which consistent with the original image features, thus limiting the powerful generative capacity of the diffusion model, and vice versa. As shown in Fig. 5, we show visual comparisons to demonstrate the effectiveness of using both FR and NR metrics. More IQA-metrics analysis about mutual interaction within FR or NR metrics can be found in the supplementary materials.

Effect of the progressive feature injection. To demonstrate the effectiveness of the PFJ module, we adopt several different configurations to validate the influence of the PFJ module. CC and CA indicate concat and cross-attention mechanisms, respectively. As shown in Tab. 5, when original image features are fused in the concat mechanism, it leads to significant performance degradation. With the cross attention mechanism, the reconstruction performance improves obviously. Furthermore, the PFJ module uses modulation coefficients to regulate the original image features, and then injects these features into the denoising process, leading to a better trade-off between fidelity and realism.

Effect between different region slicing strategy. As mentioned in Sec. 3.4, we slice the input LR images into multiple overlapping regions with the fixed resolution as the pre-trained DMs. However, the above slicing strategy is not consistent with the original intent of distinguishing between foreground and background regions as mentioned in Fig. 1. To demonstrate the effectiveness of the above slicing strategy, we apply super-pixel segmentation methods [21,68,69,75] to distinguish the foreground and background regions, and the experimental results are shown in Tab. 6, the different slicing strategy has no significant impact on the visual results. However, due to the input resolution limitation of pre-trained diffusion models, we slice or fill these segmentation results into fixed resolution regions using zero-padding, then feed these regions into the proposed AdaDiffSR for inference, which may lead to the meaningless filled regions consuming massive computational resources, thus slowing the inference time. As shown in Fig. 6, we visualize the effectiveness of the region integration strategy. With this strategy, the discontinuities disappeared significantly.

14 Y. Fan et al.

Table 6: Ablation studies of region slicing strategy on DIV2K Valid dataset [47]. We report the inference time on the 512×512 resolution.

Exp.	#P(M)	$PSNR\uparrow$	$SSIM^{\uparrow}$	LPIPS↓	$\operatorname{RT}(s)$
SAM [21]	636	24.21	0.7368	0.2232	9.43
FastSAM [75]	68	24.11	0.7323	0.2182	8.79
MobileSAM [69]	10	24.14	0.7363	0.2178	<u>8.59</u>
EfficientSAM [68]	25	24.36	0.7311	0.2204	8.72
Ours	0	24.25	0.7355	0.2153	4.81

Table 7: Ablation studies of different modules in AdaDiffSR on DIV2K Valid dataset [47]. We report the inference time on the 1024×1024 resolution.

Exp.	Models.	#P(M)	$PSNR\uparrow$	$\rm SSIM\uparrow$	LPIPS↓	RT(s)
Baseline	LDM [41]	113.60	20.58	0.5590	0.2556	22.3
(-)	+ DTSS (IQA)	113.60	22.16	0.5377	0.2843	21.5
(a)	+ DTSS (MMLE)	115.79	22.38	0.5294	0.2715	18.1
(b)	+ PFJ	120.43	<u>24.22</u>	0.7299	0.2338	16.5
(c)	+ RIS (Ours)	120.43	24.25	0.7355	0.2153	<u>16.8</u>

Effect of different modules in AdaDiffSR. To demonstrate the effectiveness of components in AdaDiffSR, we add each component to the baseline model LDM, and the results are shown in Tab. 7. DTSS (IQA) and DTSS (MMLE) denote that calculate the IQA metrics using PYIQA toolbox [5] (*i.e.*, the groundtruth) and using regressor, respectively. There is no significant gap between the above strategy, but the MMLE regressor would reduce the inference time considerably. Besides, with the addition of the PFJ module, the experimental results and inference time become better. Although the region integration strategy (RIS) did not contribute significantly to the metrics, as shown in Fig. 6, this strategy eliminated the discontinuity, thus improving the visual quality.

5 Limitations

Although the proposed AdaDiffSR achieves the trade-off between computational resources and restoration quality compared to other DMs-based SR methods. However, the DMs-based SR approaches require multiple denoising steps to obtain the restored results, which is slower than previous CNN or GAN-based SR methods. Moreover, since the fixed input resolution limitations of the pre-trained DMs, the proposed method applies the static slicing strategy, which is somewhat different from the original design intention mentioned in Fig. 1. Our future works will concentrate on investigating new mechanisms that acclerate the denoising process, and achieving more fine-grained region reconstruction.

6 Conclusion

In this paper, we explore the trends that apply adaptive inference strategy on diffusion models based SR methods. To this end, we propose a novel DMs-based SR pipeline AdaDiffSR for real-world image super-resolution, which utilizes the latent information gain to adjust timesteps dynamically during the denoising process and effectively reduces the computational resource overhead. Experimental results on both synthetic and real-world benchmarks demonstrate that the proposed AdaDiffSR achieves a better trade-off between computational resources and restoration quality. We believe that our exploration will provide more inspiration for future works.

15

Acknowledgements

This work was supported in part by the NSFC under Grant 62272380 and 62103317, the Fundamental Research Funds for the Central Universities, China (xzy022023051), the Innovative Leading Talents Scholarship of Xi'an Jiaotong University.

References

- 1. Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In: CVPRW (July 2017)
- Blau, Y., Mechrez, R., Timofte, R., Michaeli, T., Zelnik-Manor, L.: The 2018 pirm challenge on perceptual image super-resolution. In: ECCVW. pp. 0–0 (2018)
- 3. Cai, Jianrui, Z., etal.: Toward real-world single image super-resolution: A new benchmark and a new model. In: ICCV (2019)
- Che, T., Li, Y., Jacob, A.P., Bengio, Y., Li, W.: Mode regularized generative adversarial networks. arXiv preprint arXiv:1612.02136 (2016)
- 5. Chen, C., Mo, J.: IQA-PyTorch: Pytorch toolbox for image quality assessment
- Chen, C., Mo, J., Hou, J., Wu, H., Liao, L., Sun, W., Yan, Q., Lin, W.: Topiq: A top-down approach from semantics to distortions for image quality assessment. arXiv preprint arXiv:2308.03060 (2023)
- 7. Chen, C., Shi: Real-world blind super-resolution via feature matching with implicit high-resolution priors. In: ACM MM (2022)
- Chen, C., Zhou, S., Liao, L., Wu, H., Sun, W., Yan, Q., Lin, W.: Iterative token evaluation and refinement for real-world super-resolution. In: AAAI. vol. 38, pp. 1010–1018 (2024)
- 9. Ding, K., Ma, K., Wang, S., Simoncelli, E.P.: Image quality assessment: Unifying structure and texture similarity. IEEE TPAMI (2020)
- Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. IEEE TPAMI (2015)
- 11. Dong, C., Loy, C.C., Tang, X.: Accelerating the super-resolution convolutional neural network. In: ECCV. Springer (2016)
- 12. Fritsche, M., Gu, S., Timofte, R.: Frequency separation for real-world superresolution. In: ICCVW. IEEE (2019)
- 13. Golestaneh, S.A., Dadsetan: No-reference image quality assessment via transformers, relative ranking, and self-consistency. In: WACV
- Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control. arXiv preprint arXiv:2208.01626 (2022)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. NeurIPS 33 (2020)
- Ignatov, A., Kobyshev: Dslr-quality photos on mobile devices with deep convolutional networks. In: ICCV (2017)
- 17. Ji, X., Cao, Y., Tai, Y., Wang, C., Li, J., Huang, F.: Real-world super-resolution via kernel estimation and noise injection. In: CVPRW (June 2020)
- Jiménez, A.B.: Mixture of diffusers for scene composition and high resolution image generation. arXiv preprint arXiv:2302.02412 (2023)
- 19. Kang, L., Ye, P., Li, Y., Doermann, D.: Convolutional neural networks for noreference image quality assessment. In: CVPR. pp. 1733–1740 (2014)

- 16 Y. Fan et al.
- Ke, J., Wang, Q., Wang, Y., Milanfar, P., Yang, F.: Musiq: Multi-scale image quality transformer. In: ICCV (2021)
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. arXiv:2304.02643 (2023)
- 22. Kong, X., Zhao, H., Qiao, Y., Dong, C.: ClassSR: A general framework to accelerate super-resolution networks by data characteristic. In: CVPR (2021)
- Lao, S., Gong, Y., Shi, S., Yang, S., Wu, T., Wang, J., Xia, W., Yang, Y.: Attentions help cnns see better: Attention-based hybrid image quality assessment network. In: CVPRW (2022)
- 24. Ledig, C., Theis, L., Huszár, F., etal.: Photo-realistic single image super-resolution using a generative adversarial network. In: CVPR (2017)
- Liang, J., Zeng, H., Zhang, L.: Efficient and degradation-adaptive network for realworld image super-resolution. In: ECCV. Springer (2022)
- Liebel, L., Körner, M.: Single-image super resolution for multispectral remote sensing data using convolutional neural networks. ISPRS (2016)
- 27. Lim, B., Son, S., Kim, H., Nah, S., Mu Lee, K.: Enhanced deep residual networks for single image super-resolution. In: CVPRW (2017)
- Lin, X., He, J., Chen, Z., Lyu, Z., Fei, B., Dai, B., Ouyang, W., Qiao, Y., Dong, C.: Diffbir: Towards blind image restoration with generative diffusion prior. arXiv preprint arXiv:2308.15070 (2023)
- Liu, C., Wang, X., Fan, Y., Li, S., Qian, X.: Decoupling degradations with recurrent network for video restoration in under-display camera. In: AAAI. vol. 38, pp. 3558– 3566 (2024)
- Liu, C., Wang, X., Xu, X., Tian, R., Li, S., Qian, X., Yang, M.H.: Motion-adaptive separable collaborative filters for blind motion deblurring. In: CVPR. pp. 25595– 25605 (2024)
- Liu, C., Yang, H., Fu, J., Qian, X.: Learning trajectory-aware transformer for video super-resolution. In: CVPR. pp. 5687–5696 (2022)
- 32. Liu, M., Zhang, Z., Hou, L., Zuo, W., Zhang, L.: Deep adaptive inference networks for single image super-resolution. In: ECCVW. pp. 131–148. Springer (2020)
- Ma, C., Yang, C.Y., Yang, X., Yang, M.H.: Learning a no-reference quality metric for single-image super-resolution. CVIU (2017)
- Maeda, S.: Unpaired image super-resolution using pseudo-supervision. In: CVPR (2020)
- Mao, Q., Lee, H.Y., Tseng, H.Y., Ma, S., Yang, M.H.: Mode seeking generative adversarial networks for diverse image synthesis. In: CVPR (2019)
- Mei, K., Delbracio, M., Talebi, H., Tu, Z., Patel, V.M., Milanfar, P.: Codi: Conditional diffusion distillation for higher-fidelity and faster image generation. In: CVPR. pp. 9048–9058 (2024)
- Mittal, A., Moorthy, A.K., Bovik, A.C.: Blind/referenceless image spatial quality evaluator. In: ASILOMAR. pp. 723–727 (2011)
- Mittal, A., Moorthy, A.K., Bovik, A.C.: No-reference image quality assessment in the spatial domain. IEEE TIP 21(12), 4695–4708 (2012)
- Mittal, A., Soundararajan, R., Bovik, A.C.: Making a "completely blind" image quality analyzer. SPL 20(3), 209–212 (2012)
- Mokady, R., Hertz, A., Aberman: Null-text inversion for editing real images using guided diffusion models. In: CVPR. pp. 6038–6047 (2023)
- 41. Rombach, R., Blattmann, A., etal.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)

17

- 42. Sahak, H., Watson, D., Saharia, C., Fleet, D.: Denoising diffusion probabilistic models for robust image super-resolution in the wild. arXiv preprint arXiv:2302.07864 (2023)
- Saharia, C., Ho, J., etal.: Image super-resolution via iterative refinement. IEEE TPAMI (2022)
- 44. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
- 45. Srivastava, A., Valkov, L., Russell, C., Gutmann, M.U., Sutton, C.: Veegan: Reducing mode collapse in gans using implicit variational learning. NeurIPS (2017)
- 46. Su, S., Yan, Q., Zhu, Y., Zhang, C., Ge, X., Sun, J., Zhang, Y.: Blindly assess image quality in the wild guided by a self-adaptive hyper network. In: CVPR (2020)
- Timofte, R., Agustsson: Ntire 2017 challenge on single image super-resolution: Methods and results. In: CVPRW (2017)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. NeurIPS **30** (2017)
- Verelst, T., Tuytelaars, T.: Dynamic convolutions: Exploiting spatial sparsity for faster inference. In: CVPR. pp. 2320–2329 (2020)
- Wang, J., Chan, K.C., Loy, C.C.: Exploring clip for assessing the look and feel of images. In: AAAI (2023)
- 51. Wang, J., Yue, Z., Zhou.: Exploiting diffusion prior for real-world image superresolution. arXiv preprint arXiv:2305.07015 (2023)
- Wang, L., Dong, X., Wang, Y., Ying, X., Lin, Z., An, W., Guo, Y.: Exploring sparsity in image super-resolution for efficient inference. In: CVPR. pp. 4917–4926 (2021)
- 53. Wang, L., Wang, Y., Dong, X., Xu, Q., Yang, J., An, W., Guo, Y.: Unsupervised degradation representation learning for blind super-resolution. In: CVPR (2021)
- 54. Wang, S., Liu, J., etal.: Adaptive patch exiting for scalable single image superresolution. In: ECCV. Springer (2022)
- 55. Wang, X., Xie, L., etal.: Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In: ICCV (2021)
- Wang, X., Yu, K., Dong, C., Loy, C.C.: Recovering realistic texture in image superresolution by deep spatial feature transform. In: CVPR (2018)
- 57. Wang, X., Yu, K., Wu, S., etal.: ESRGAN: Enhanced super-resolution generative adversarial networks. In: ECCVW (2018)
- 58. Wei, P., Xie, Z., Lu, H., etal.: Component divide-and-conquer for real-world image super-resolution. In: ECCV. Springer (2020)
- Wu, J.Z., Ge, Y., Wang, X., Lei, S.W., Gu, Y., Shi, Y., Hsu, W., Shan, Y., Qie, X., Shou, M.Z.: Tune-a-video: One-shot tuning of image diffusion models for textto-video generation. In: ICCV (2023)
- Wu, R., Yang, T., Sun, L., Zhang, Z., Li, S., Zhang, L.: Seesr: Towards semanticsaware real-world image super-resolution. In: CVPR. pp. 25456–25467 (2024)
- Xue, W., Zhang, L., etal.: Gradient magnitude similarity deviation: A highly efficient perceptual image quality index. IEEE TIP (2013)
- Yang, S., Wu, T., Shi, S., Lao, S., Gong, Y., Cao, M., Wang, J., Yang, Y.: Maniqa: Multi-dimension attention network for no-reference image quality assessment. In: CVPRW (2022)
- 63. Ying, Z., Niu, H., Gupta, P., Mahajan, D., Ghadiyaram, D., Bovik, A.: From patches to pictures (paq-2-piq): Mapping the perceptual space of picture quality. In: CVPR. pp. 3575–3585 (2020)
- 64. Yu, J., Huang, T.: Autoslim: Towards one-shot architecture search for channel numbers. arXiv preprint arXiv:1903.11728 (2019)

- 18 Y. Fan et al.
- 65. Yu, J., Huang, T.S.: Universally slimmable networks and improved training techniques. In: ICCV (2019)
- Yu, J., Yang, L., Xu, N., Yang, J., Huang, T.: Slimmable neural networks. arXiv preprint arXiv:1812.08928 (2018)
- 67. Yue, Z., Wang, J., Loy, C.C.: Resshift: Efficient diffusion model for image superresolution by residual shifting. NeurIPS (2024)
- Yunyang Xiong, Bala Varadarajan, L.W., etal.: Efficientsam: Leveraged masked image pretraining for efficient segment anything. arXiv:2312.00863 (2023)
- Zhang, C., Han, D., Qiao, Y., Kim, J.U., Bae, S.H., Lee, S., Hong, C.S.: Faster segment anything: Towards lightweight sam for mobile applications (2023)
- 70. Zhang, J., Lu, S., Zhan, F., Yu, Y.: Blind image super-resolution via contrastive representation learning. arXiv preprint arXiv:2107.00708 (2021)
- Zhang, K., Liang, J., Van Gool, L., Timofte, R.: Designing a practical degradation model for deep blind image super-resolution. In: ICCV (2021)
- Zhang, L., Zhang, L., Bovik, A.C.: A feature-enriched completely blind image quality evaluator. IEEE TIP 24(8), 2579–2591 (2015)
- 73. Zhang, R., Isola, P., Efros, etal.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)
- 74. Zhang, W., Ma, K., Yan, J., Deng, D., Wang, Z.: Blind image quality assessment using a deep bilinear convolutional neural network. IEEE TCSVT (2020)
- 75. Zhao, X., Ding, W., An, Y., Du, Y., Yu, T., Li, M., Tang, M., Wang, J.: Fast segment anything (2023)
- Zheng, H., Yang, H., Fu, J., Zha, Z.J., Luo, J.: Learning conditional knowledge distillation for degraded-reference image quality assessment. In: ICCV (2021)