

DiT4SR: Taming Diffusion Transformer for Real-World Image Super-Resolution

Zheng-Peng Duan^{1,2} * Jiawei Zhang² Xin Jin¹ Ziheng Zhang¹ Zheng Xiong²

Dongqing Zou^{2,3} Jimmy S. Ren^{2,4} Chunle Guo^{1,5} Chongyi Li^{1,5} †

¹VCIP, CS, Nankai University ²SenseTime Research ³PBVR

⁴Hong Kong Metropolitan University ⁵NKIARI, Shenzhen Futian

<https://adam-duan.github.io/projects/dit4sr/>



Figure 1. Performance comparison between our DiT4SR and other state-of-the-art methods on two real-world LR images. Thanks to the powerful generative capabilities of SD3 and our delicate designs, which allow sufficient interaction between the input LR information and the generation process, our method produces more realistic details while maintaining fidelity, especially in text, structures, and details.

Abstract

Large-scale pre-trained diffusion models are becoming increasingly popular in solving the Real-World Image Super-Resolution (Real-ISR) problem because of their rich generative priors. The recent development of diffusion transformer (DiT) has witnessed overwhelming performance over the traditional UNet-based architecture in image gen-

eration, which also raises the question: Can we adopt the advanced DiT-based diffusion model for Real-ISR? To this end, we propose our DiT4SR, one of the pioneering works to tame the large-scale DiT model for Real-ISR. Instead of directly injecting embeddings extracted from low-resolution (LR) images like ControlNet, we integrate the LR embeddings into the original attention mechanism of DiT, allowing for the bidirectional flow of information between the LR latent and the generated latent. The sufficient interaction of these two streams allows the LR stream to evolve with the

*This project is done during the internship at SenseTime Research.

†Corresponding author.

diffusion process, producing progressively refined guidance that better aligns with the generated latent at each diffusion step. Additionally, the LR guidance is injected into the generated latent via a cross-stream convolution layer, compensating for DiT’s limited ability to capture local information. These simple but effective designs endow the DiT model with superior performance in Real-ISR, which is demonstrated by extensive experiments.

1. Introduction

Real-ISR [55, 69] aims to recover a high-resolution (HR) image from its LR version under various types of degradations, such as compression, blur, and noise. Unlike traditional ISR [19], Real-ISR requires that the model can not only remove the complex degradations, but also generate perceptually realistic details to enhance the visual quality. The high ill-posedness of this task places demands on prior knowledge of the model. Without such prior, the model is prone to generating ambiguous results and suffers from degraded restoration quality. Therefore, researchers have turned their attention to large-scale pre-trained text-to-image (T2I) models [11, 44], especially Stable Diffusion (SD), which are trained on billions of high-quality nature images and encompass rich real-world prior knowledge. They build their methods upon UNet-based SD, known as SD1 [43], SD2 [43], and SDXL [39], and treat input LR images as conditions via ControlNet [70] or ControlNet-like manners to generate the corresponding HR images.

Recently, the development of diffusion transformer (DiT) [38] has made DiT-based architecture prevalent. Owing to the overwhelming performance of SD3 [21] and Flux [3], MM-DiT has been proven to be an effective transformer block for generative models. It features two learnable streams for visual features and text tokens and enables a bidirectional flow of information between the two modalities through attention operation. Building on such architectural innovations, along with other techniques like scaling up, DiT-based diffusion models demonstrate promising performance in detail generation and image quality. *Given these merits, a pivotal question arises: Can we adopt the advanced DiT-based diffusion model for Real-ISR?*

An intuitive solution is to adopt ControlNet [70] as the DiT controller and inject LR information to guide the generative process for Real-ISR, which is shown in Figure 2 (a)¹. Specifically, several MM-DiT blocks are duplicated as a trainable copy to initialize ControlNet. The LR latent, obtained by passing the LR image through the pre-trained VAE encoder, is used as input. The hidden state output from each block in ControlNet is then directly added to the Noise

¹The implementation of SD3-ControlNet follows the default code and configuration provided by Diffusers [50].

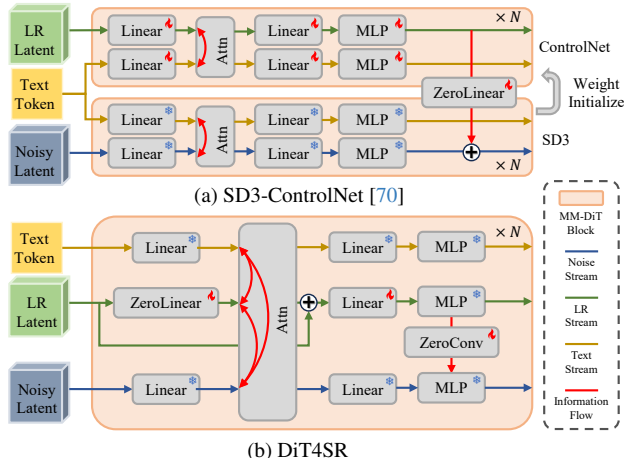


Figure 2. Network Structure comparison between SD3-ControlNet and our DiT4SR. The information flow across streams is marked with red lines and arrows indicate the direction. Notably, our DiT4SR enables bidirectional information interaction, allowing the LR Stream and Noise Stream to continuously interact and evolve together, whereas SD3-ControlNet relies on one-directional information flow, restricting the interaction.

Stream of SD3 through a trainable convolution layer. Since ControlNet is originally designed for UNet-based architectures, directly adding LR information ignores the unique characteristics of DiT, restricting information interactions and potentially limiting performance.

To leverage the strength of DiT, we abandon the traditional technical roadmap such as ControlNet and propose a new control architecture designed specifically for DiT, named DiT4SR. Instead of processing the LR information in the additionally created DiT blocks like Figure 2 (a), we integrate the LR Stream into the DiT blocks, allowing for more efficient interaction between LR information and the diffusion process. Considering that both the LR latent and the noisy latent are forms of visual information, we adopt a similar design for the LR Stream as the Noise Stream. In each MM-DiT block, our method duplicates the modules for the Noise Stream to process the LR input. Such **LR Integration in Attention** allows for the bidirectional information interaction between the LR Stream and the Noise Stream, which is shown in Figure 2 (b).

The bidirectional interaction between the two streams allows the LR Stream to evolve alongside the diffusion process, enabling it to generate progressively refined and context-aware guidance that better aligns with the generated latent. To preserve LR information and enhance consistency in deeper blocks, an additional **LR Residual** is introduced to directly connect the input and output of the attention module within the LR Stream. However, since the attention mechanism operates globally, it alone is not sufficient for SR tasks. Capturing local information is equally essential to restore fine details. Therefore, we further inject the LR guidance into the Noise Stream through the convo-

lution layer. The **LR Injection between MLP** allows our model to aggregate the local information, compensating for the limited local information-capturing ability of DiT [61].

Our **contributions** can be summarized as follows:

- To our knowledge, our method is one of the pioneering works to tame the large-scale DiT model for Real-ISR.
- Instead of copying blocks like ControlNet, we integrate the LR Stream into the original DiT block, which enables bidirectional information interaction between the LR guidance and the diffusion process.
- We introduce a convolutional layer for injecting LR guidance into the Noise Stream, which compensates for the limited local information-capturing ability of DiT.

2. Related Work

Deep learning-based ISR methods have achieved significant progress, where the architectures develop from convolutional networks [18, 19, 34, 73–75] to transformers [10, 13, 14, 31]. However, these methods struggle to handle Real-ISR due to the complex degradations in real-world scenarios and the ill-posed nature of this task. To address the former issue, BSRGAN [69] and Real-ESRGAN [55] explore more complex degradation models. For the ill-posed property, several GAN-based methods [6, 9, 33, 54] are proposed to learn the distribution of high-quality images and generate perceptually realistic details. However, these methods still suffer from training instability and generate unnatural visual artifacts [8, 32, 62].

With the success of large-scale T2I models such as SD [43], researchers [17, 22, 28, 40, 45, 46, 49, 51, 67] have shifted their attention to use the generative priors embedded in these powerful pre-trained models to handle the Real-ISR problem. Built upon the first-generation SD [43], which is based on a UNet architecture, StableSR [53] and DiffBIR [35] inject the LR information as conditions via ControlNet [70] or ControlNet-like manners. PASD [65] and SeeSR [60] further incorporate high-level semantic information to guide the diffusion process. Employing the more advanced SDXL, SUPIR and FaithDiff [12] investigate the effects of scaling up in Real-ISR. There are also some other methods [7, 29, 36, 37, 47, 56, 59, 63, 68] committed to a more efficient diffusion process. With the development of DiT, DiT-SR [15] trains DiT-based SR models from scratch. DreamClear [2] and DPIR [27] propose DiT-based image restoration models, but still adopt the ControlNet to inject LR information, which prevents it from fully leveraging the advantage of DiT. From the perspective of network architecture, our DiT4SR is one of the pioneering works to tame the large-scale diffusion transformer model for Real-ISR, while discarding the ControlNet-like approach. There are two concurrent methods [20, 30] that explore one-step DiT-based SR models in terms of flow trajectory.

3. Methodology

Given a LR image I_{LR} undergoing complex degradations, Real-ISR aims to recover the corresponding HR image R . Along with the removal of degradations, Real-ISR simultaneously requires the model to generate realistic details, thereby enhancing visual perception. To overcome the ill-posedness of this task, researchers build their methods on large-scale pre-trained T2I models, especially SD, to leverage the rich generative priors. The recent development of DiT has gradually replaced the traditional UNet-based diffusion model, achieving state-of-the-art performance in detail generation and visual quality. In this study, we focus on leveraging the strength of DiT and tame it for Real-ISR.

To this end, we propose our DiT4SR, which is built on DiT-based SD3. The overview of the architecture of our DiT4SR is introduced in Section 3.1. Instead of creating additional DiT blocks to process the LR information like SD3-ControlNet, we integrate the LR Stream into the original attention calculation of DiT and allow bidirectional interaction of information with the Noise Stream and Text Stream, which is detailed in Section 3.2. However, relying solely on the attention mechanism is not enough to make the DiT competent for Real-ISR. To inject more LR guidance and capture the local information of LR features, the intermediate features of MLP in the LR Stream are introduced into that of the Noise Stream through a depth-wise convolution layer, which are detailed in Section 3.3.

3.1. Overview of Architecture

Since our DiT4SR is built on the popular DiT-architected SD3, we first introduce SD3 briefly. Like previous SD models [43], SD3 also conducts the diffusion process in the latent space. It consists of a sequence of MM-DiT blocks, where separate weight sets handle text and image embeddings, forming the Noise Stream and the Text Stream, as shown in the bottom part of Figure 2 (a). This structure employs two independent transformers for each modality while merging their sequences using the attention mechanism, enabling cross-modal interaction. The bidirectional interaction enables both streams to evolve together throughout the diffusion process, representing the key strength of the DiT. If ControlNet [50, 70] is adopted as the DiT controller for Real-ISR, as illustrated in Figure 2 (a), SD3-ControlNet processes the LR Stream within additional DiT blocks and directly injects the LR embedding into the Noise Stream via trainable linear layers. This approach establishes only a one-way information flow from the LR Stream to the Noise Stream, limiting information interaction. In contrast, our DiT4SR integrates the LR Stream directly into the original DiT blocks, as illustrated in Figure 2 (b). This design enables bidirectional information flow, allowing the LR Stream to continuously adapt throughout the diffusion process and generate guidance that aligns more effectively with




Figure 3. Overview of DiT4SR (a) and MM-DiT-Control block (b). Our DiT4SR integrates LR Stream to MM-DiT-Control Block. **LR Integration in Attention** allows the bidirectional interaction between LR Stream and Noise Stream. **LR Residual** is introduced to enhance the consistency of LR guidance. Additionally, **LR Injection between MLP** injects the LR guidance from LR Stream to Noise Stream through a trainable convolution layer.

the evolving Noise Stream.

As illustrated in Figure 3, our DiT4SR shares a similar architecture with SD3. Specifically, before being fed into the diffusion transformer, the noisy latent $Z \in \mathbb{R}^{H \times W \times C}$ is first flattened into a patch sequence of length K , where $K = \frac{H}{2} \cdot \frac{W}{2}$, and then projected to a D -dimensional space using a linear layer. Then, the position embedding is added to the obtained noisy image token $X \in \mathbb{R}^{K \times D}$. To introduce the LR information, we first encode the LR image I_{LR} into the latent space through the pre-trained VAE encoder. Considering the LR latent and noisy latent are both forms of visual representations, we follow the same process and add the same position embedding to obtain the LR image token $L \in \mathbb{R}^{K \times D}$. Following SD3 [21], the input caption that describes I_{LR} is encoded by three pre-trained text models, including CLIP-L [41], CLIP-G [16], and T5 XXL [42]. The text representations output by two CLIP models are pooled and combined with the timestep t to modulate the internal features of DiT. Besides the pooled text representations, the text token $C \in \mathbb{R}^{M \times D}$ is also constructed with a length of M by combining all three text representations.

With the noisy image token X and the text token C as input, the original MM-DiT block adopts two separate sets of

weights for these two modalities, named Noise Stream and Text Stream. In our DiT4SR, we introduce an additional stream, named LR Stream, to tackle the LR image token L . The MM-DiT block is modified as the MM-DiT-Control block, which allows the LR information to guide the generation of HR latent. After N MM-DiT-Control blocks and the unpatch operation, Noise Stream outputs the denoised latent for this timestep t . Repeating the diffusion process for T steps and decoding the clean latent Z_0 , we can finally obtain the desired HR result R . In the following two sections, the detailed modifications in the MM-DiT-Control block are introduced.

3.2. LR Integration in Attention

Following the design of the original MM-DiT, all three streams sequentially undergo the joint attention mechanism and the MLP operation in the MM-DiT-Control block. As for the joint attention, the input can be formulated as

$$\begin{aligned}
 Q &= P_Q^X(X) \odot P_Q^L(L) \odot P_Q^C(C), \\
 K &= P_K^X(X) \odot P_K^L(L) \odot P_K^C(C), \\
 V &= P_V^X(X) \odot P_V^L(L) \odot P_V^C(C),
 \end{aligned} \tag{1}$$

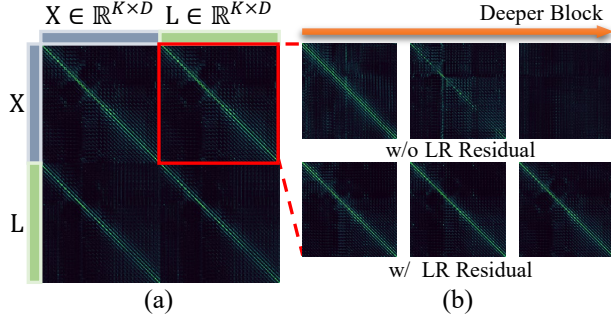


Figure 4. (a) Visualization of four attention maps for noisy image token \mathbf{X} and LR image token \mathbf{L} ($\mathbf{X} \rightarrow \mathbf{X}$, $\mathbf{X} \rightarrow \mathbf{L}$, $\mathbf{L} \rightarrow \mathbf{X}$, $\mathbf{L} \rightarrow \mathbf{L}$) in the 24th MM-DiT-Control. Diagonal lines of $\mathbf{X} \rightarrow \mathbf{L}$ and $\mathbf{L} \rightarrow \mathbf{X}$ indicate the information interaction between \mathbf{X} and \mathbf{L} . (b) Attention maps for $\mathbf{X} \rightarrow \mathbf{L}$ w/ and w/o LR Residual. Without LR Residual, LR guidance diminishes with increasing block depth (in the 1st, 13th, and 24th MM-DiT-Control). LR Residual can significantly enhance the consistency of LR guidance.

where $P_Q^X, P_K^X, P_V^X, P_Q^L, P_K^L, P_V^L$ are pre-trained fixed linear projections for \mathbf{X} and \mathbf{L} . \odot denotes the operation of concatenation in the token length dimension. P_Q^L, P_K^L , and P_V^L are newly created trainable linear projections for \mathbf{L} , whose weights are initialized to zeros. In this way, the impact of \mathbf{L} can be ignored at the beginning of the training, and progressively grows during the training. The joint attention in MM-DiT-Control is calculated as

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\underbrace{\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}}}_{\text{attention map}}\right) \mathbf{V}, \quad (2)$$

which allows for more comprehensive interaction between these three streams. We further visualize the attention maps regarding LR image token \mathbf{X} and noisy image token \mathbf{L} in Figure 4 (a). We observe that the diagonals are activated obviously in both self-attention regions and cross-attention regions, which indicates the information interaction between the corresponding positions on \mathbf{X} and \mathbf{L} . This bidirectional interaction not only enables the Noise Stream to be influenced by LR guidance but also allows the LR Stream to adapt based on the state of the Noise Stream, providing more accurate and context-aware guidance.

Another noteworthy thing is the information interaction between \mathbf{L} and \mathbf{X} decay through successive attention blocks, which is shown in Figure 4 (b). The possible reason is that, under the influence of joint attention, the LR information evolves alongside the Noise Stream and may experience undesired disruptions. This gradual degradation reduces its effectiveness in guiding the Noise Stream. To enhance the consistency of LR guidance, an additional shortcut directly transfers the input LR information to the output of the joint attention mechanism. By incorporating the **LR Residual**, the LR guidance is effectively preserved




Figure 5. (a) is the LR input. (b) is the result w/o LR Injection between MLP. (c) injects the LR information through a linear layer. (d) is the result of our DiT4SR, which replaces the linear layer with a convolution layer. The convolution layer helps capture more precise local information, reflected by remarkable performance in recovering fine structures. All results are obtained by retraining the specific models under the same settings.

in deeper transformer blocks, ensuring its consistent influence on the Noise Stream throughout the diffusion process. The visual comparison between the models with and without LR Residual is presented in Figure 7 (c) and (f).

3.3. LR Injection between MLP

Since joint attention operates at a global level and relies solely on position embeddings to provide spatial information, it is insufficient for effectively adapting DiT to Real-ISR. This limitation arises because local information is equally important for restoring fine details, which global attention alone may overlook. As shown in Figure 5 (b), relying solely on the joint attention still faces a challenge in restoring the text and maintaining image fidelity.

To strengthen the guidance of LR information, we further connect the intermediate features of the MLPs from LR Stream to Noisy Stream with a 3×3 depth-wise convolution layer whose weights are initialized to zeros. Specifically, in both MLPs for the LR Stream and Noisy Stream, the hidden state dimensions are first expanded by a factor of 4 and then projected back to the original size through two linear projections. We denote these two intermediate features as $\phi(\mathbf{X}) \in \mathbb{R}^{K \times 4D}$ and $\eta(\mathbf{L}) \in \mathbb{R}^{K \times 4D}$. We first reshape $\eta(\mathbf{L})$ to $\eta(\mathbf{L})' \in \mathbb{R}^{\frac{K}{2} \times \frac{W}{2} \times 4D}$, and then pass it through the 3×3 depth-wise convolution layer. After reshaping back to the image token form, the LR information is effectively injected into the Noise Stream.

Meanwhile, the 3×3 depth-wise convolution layer also compensates for the limited local information-capturing ability of DiT. As mentioned, the joint attention is calculated at a global level, thus lacking the local guidance of LR information. Even if using one linear layer to inject LR guidance, it still remains challenging to repair fine structures such as letters, as shown in Figure 5 (c). The 3×3 depth-wise convolution helps capture more precise local information from the LR Stream and inject such LR guidance into the Noise Stream, achieving better performance, especially for fine structures which are shown in Figure 5 (d).

Datasets	Metrics	Real-ESRGAN	SwinIR	ResShift	StableSR	SeeSR	DiffBIR	OSEDiff	SUPIR	DreamClear	SD3-ControlNet	DiT4SR
DrealSR	LPIPS ↓	0.282	0.274	0.353	0.273	0.317	0.452	0.297	0.419	0.354	0.323	0.365
	MUSIQ ↑	54.267	52.737	52.392	58.512	65.077	65.665	64.692	59.744	44.047	55.956	64.950
	MANIQA ↑	0.490	0.475	0.476	0.559	0.605	0.629	0.590	0.552	0.455	0.545	0.627
	ClipIQA ↑	0.409	0.396	0.379	0.438	0.543	0.572	0.519	0.518	0.379	0.449	0.548
	LIQE ↑	2.927	2.745	2.798	3.243	4.126	3.894	3.942	3.728	2.401	3.059	3.964
RealSR	LPIPS ↓	0.271	0.254	0.316	0.306	0.299	0.347	0.292	0.357	0.325	0.305	0.319
	MUSIQ ↑	60.370	58.694	56.892	65.653	69.675	68.340	69.087	61.929	59.396	62.604	68.073
	MANIQA ↑	0.551	0.524	0.511	0.622	0.643	0.653	0.634	0.574	0.546	0.599	0.661
	ClipIQA ↑	0.432	0.422	0.407	0.472	0.577	0.586	0.552	0.543	0.474	0.484	0.550
	LIQE ↑	3.358	2.956	2.853	3.750	4.123	4.026	4.065	3.780	3.221	3.338	3.977
RealLR200	MUSIQ ↑	62.961	63.548	59.695	63.433	69.428	68.027	69.547	64.837	65.926	65.623	70.469
	MANIQA ↑	0.553	0.560	0.525	0.579	0.612	0.629	0.606	0.600	0.597	0.587	0.645
	ClipIQA ↑	0.451	0.463	0.452	0.458	0.566	0.582	0.551	0.524	0.546	0.526	0.588
	LIQE ↑	3.484	3.465	3.054	3.379	4.006	4.003	4.069	3.626	3.775	3.733	4.331
RealLQ250	MUSIQ ↑	62.514	63.371	59.337	56.858	70.556	69.876	69.580	66.016	66.693	66.385	71.832
	MANIQA ↑	0.524	0.534	0.500	0.504	0.594	0.624	0.578	0.584	0.585	0.568	0.632
	ClipIQA ↑	0.435	0.440	0.417	0.382	0.562	0.578	0.528	0.483	0.502	0.509	0.578
	LIQE ↑	3.341	3.280	2.753	2.719	4.005	4.003	3.904	3.605	3.688	3.639	4.356

Table 1. Quantitative comparison with state-of-the-art Real-ISR methods on four real-world benchmarks. Best and second best performance are highlighted in red and blue, respectively. Our DiT4SR achieves the best or comparable performance across four benchmarks.

4. Experiments

4.1. Experimental Settings

Datasets. Following SeeSR [60], we adopt a combination of images from DIV2K [1], DIV8K [23], Flickr2K [48], and the first 10K face images from FFHQ [25] during training. To fully exploit the potential of our method and scale up the training dataset, another 1K self-captured high-resolution images are included. The degradation pipeline of Real-ESRGAN [55] is utilized to synthesize LR-HR training pairs with the same parameter configuration as SeeSR. Note that the resolutions are set to 128×128 and 512×512 for LR and HR images, respectively.

Since our method specifically focuses on Real-ISR task, we evaluate our model on four widely used real-world datasets, including DrealSR [58], RealSR [5], RealLR200 [60], and RealLQ250 [2]. All experiments are conducted with the scaling factor of $\times 4$. DrealSR and RealSR respectively consist of 93 and 100 images. Following SeeSR, center-cropping is adopted for these two datasets, and the resolution of LR images is set to 128×128 . RealLR200 is proposed in SeeSR, which comprises 200 images of significantly different resolutions. RealLQ250 is established by DreamClear [2], which consists of 200 images with a fixed resolution of 256×256 . Both RealLR200 and RealLQ250 lack corresponding GT images.

Metrics. As claimed in previous studies [4, 24, 66], full-reference metrics, like PSNR and SSIM [57], are difficult to reflect the visual effects of restored results. Thus, most studies only use the perceptual measurement LPIPS [71]

for image fidelity. We additionally include an evaluation of image fidelity in the user study. We use MUSIQ [26], MANIQA [64], ClipIQA [52], and LIQE [72] as non-reference metrics to measure image quality.

Implementation is detailed in the supplementary material.

4.2. Comparison with Other Methods

We compare our method with state-of-the-art Real-ISR methods, including GAN-based methods (*i.e.* Real-ESRGAN [55] and SwinIR [31]), diffusion-based methods with UNet architecture (*i.e.* ResShift [67], StableSR [53], SeeSR [60], DiffBIR [35], OSEDiff [59], and SUPIR [66]), and diffusion-based methods with DiT architecture (*i.e.* DreamClear [2] and SD3-ControlNet). SD3-ControlNet is initialized with the SD3.5-medium parameters, employing the default configuration provided by Diffusers [50] and trained under the same settings as our approach.

Quantitative Comparisons. The quantitative comparisons with state-of-the-art Real-ISR methods on four benchmarks are presented in Table 1. For DrealSR and RealSR datasets, while SeeSR and DiffBIR demonstrate strong performance, our method achieves competitive results, closely matching or exceeding their performance. On RealLR200 and RealLQ250 datasets, our method exhibits overwhelming performance, achieving top performance across all non-reference metrics. These results highlight the capability of our method to produce high-quality restoration results.

Qualitative Comparisons. The qualitative comparisons against other methods are provided in Figure 6. From the first two rows, our method is capable of generating results

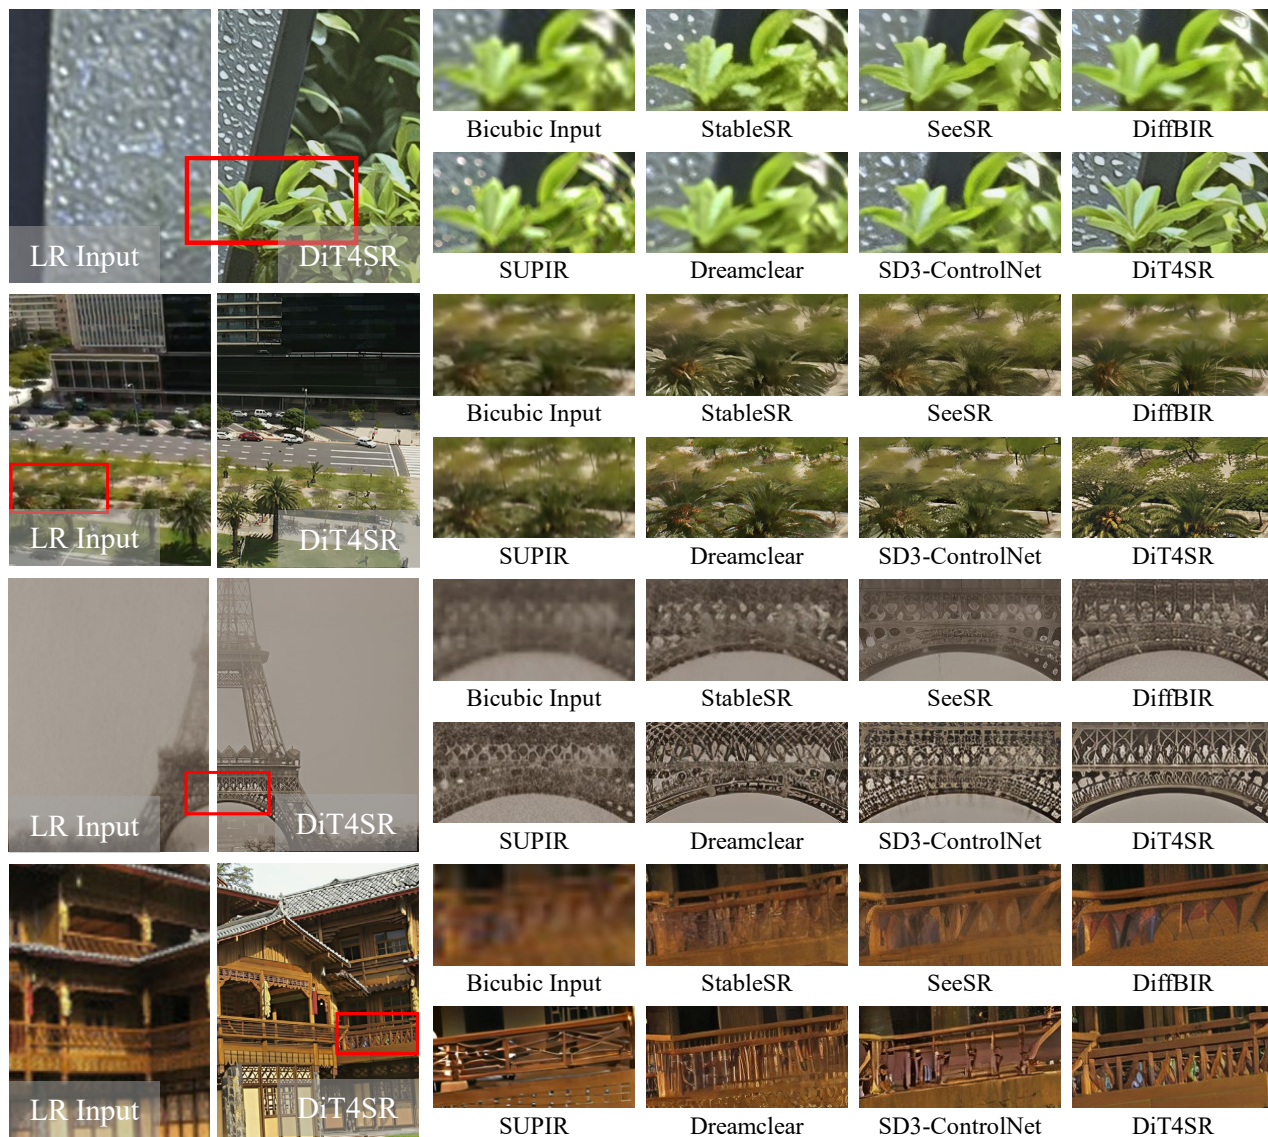


Figure 6. Qualitative comparisons with state-of-the-art Real-ISR methods on RealSR (1st row), RealLR200 (2nd and 3rd rows), and RealLO250 (4th row). Our DiT4SR achieves the best performance in terms of image realism and detail generation while maintaining fidelity to the input LR image, especially preserving fine structures. More visual results can be found in the supplementary material.

with better clarity and more abundant details over the compared methods, even when encountering severe blurring degradations. This can be attributed to our method fully leveraging the outstanding generative capabilities of SD3. Furthermore, as illustrated by the last two rows, our method exhibits a distinct advantage in processing fine structures, such as the architectural structure. Notably, even SD3-ControlNet, which is also built upon SD3, fails to handle these aspects as effectively. This further highlights the superiority of our control mechanism over ControlNet, where more comprehensive information interaction enables the model to better leverage LR information, producing high-quality restoration outcomes with promising fidelity.

User Study. To further validate the restoration quality, we invite 80 volunteers to conduct the user study. We randomly select 60 LR images from these four datasets (DrealSR, RealSR, RealLR200, and RealLQ250), and adopt the four latest methods (SeeSR, DiffBIR, SUPIR, and DreamClear) for comparison. In the user study, participants were presented with three images for each evaluation: the original LR input image, the restoration result generated by our method, and the restoration result from a randomly selected other method. Participants were asked to answer two questions: (1) Which restoration result has higher image realism? (2) Which restoration result has better fidelity to the original image content? The results reported in Table 2 demonstrate

Ours vs.	SeeSR	DiffBIR	SUPIR	DreamClear
Realism	82.1%	83.6%	81.7%	72.7%
Fidelity	68.9%	79.5%	75.4%	64.5%

Table 2. Results of user study on real-world data. The numbers indicate the winning rate of our DiT4SR over the compared method in terms of image realism and fidelity.

Model	LR Integation	LR Residual	LR Injection	MUSIQ \uparrow	MANIQA \uparrow
FULL	✓	✓	Conv	71.832	0.632
A	✗	✓	Conv	66.963	0.574
B	✓	✗	Conv	70.887	0.614
C	✓	✓	✗	71.202	0.610
D	✓	✓	Linear	71.607	0.621

Table 3. Ablation results on RealQ250 for our DiT4SR. All variants are trained using the same settings as the full model.

that our method outperforms other approaches.

5. Ablation Study

To further demonstrate the effectiveness of each component, we conduct the ablation study on RealQ250 with MUSIQ and MANIQA as evaluation metrics. All variants are trained using the same settings as the full model for fair comparisons. The detailed framework for each variant can be found in the supplementary material.

Effectiveness of LR Integration. To examine the effectiveness of LR Integration in Attention, we exclude the LR Stream from the attention computation while keeping all other components unchanged. From the results reported in Table 3, we can see that both MUSIQ and MANIQA present a significant decline, which demonstrates that relying solely on LR injection between MLP layers is insufficient to generate high-quality results. Figure 7 (b) also indicates that without bidirectional information interaction between the LR Stream and the generated latent, degradations cannot be effectively removed. This is because, without such information interaction, the LR guidance cannot adapt based on the evolving noisy latent, which limits the ability of the model to adaptively refine the restoration. Continuous interaction allows the LR information to progressively adjust and provide more accurate guidance, which is essential for effectively addressing complex degradations.

Effectiveness of LR Residual. Our intention of introducing LR residual is to maintain the consistency of LR guidance in deeper DiT blocks. As shown in Table 3, without the LR residual, both MUSIQ and MANIQA exhibit a decline in performance. From Figure 7 (c), we can see that the results contain several noticeable artifacts, which degrade the fidelity of the image content. This issue primarily stems from the unstable evolution of the LR Stream, which may




Figure 7. Visual comparison for the ablation study. Variant A, B, and C remove the LR Integation, LR Residual, and LR Injection, respectively. Variant D replaces the convolution layer with the linear layer in LR Injection.

encounter undesired disruptions. By incorporating the LR Residual, our DiT4SR effectively stabilizes the LR Stream, generating results with higher fidelity.

Effectiveness of LR Injection. First, we remove the LR Injection between MLP layers. In Table 3, we can observe a slight decline in the evaluation metrics. This indicates that relying on LR integration in attention alone can produce passable results. However, from Figure 7 (d), particularly in the eye region, we can observe noticeable content distortions. This is mainly because attention operates globally, which is not sufficient for SR tasks that also rely on local information to accurately restore fine details. We also replace the 3×3 depth-wise convolution with a linear layer. Although Table 3 shows similar performance between the linear layer and the convolution, Figure 7 (e) indicates that artifacts and distortions have not been alleviated. The 3×3 depth-wise convolution layer effectively compensates for the limited local information-capturing ability of DiT, and significantly enhances the fidelity of our DiT4SR, which cannot be fully reflected by non-reference metric data.

6. Conclusion

In this paper, we present DiT4SR, one of the pioneering works to tame the large-scale DiT model for Real-ISR. Unlike ControlNet, which directly injects embeddings extracted from LR images, our method integrates LR embeddings within the original attention mechanism of DiT. This enables bidirectional information flow between the LR latent and the generated latent. Moreover, we introduce a cross-stream convolution layer that injects LR-guided information into the generated latent. This design not only enhances LR guidance but also compensates for weak local feature-capturing ability of DiT. Through these modifications, DiT4SR achieves superior performance in Real-ISR, as demonstrated by extensive experiments. Our work highlights the potential of leveraging DiT for high-quality image restoration, paving the way for future research.

7. Acknowledgement

This work was supported in part by the National Natural Science Foundation of China (62306153, 62225604), the Natural Science Foundation of Tianjin, China (24JCJQC00020), the Young Elite Scientists Sponsorship Program by CAST (YESS20240686), the Fundamental Research Funds for the Central Universities (Nankai University, 070-63243143), and Shenzhen Science and Technology Program (JCYJ20240813114237048). The computational devices are partly supported by the Supercomputing Center of Nankai University (NKSC).

References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *CVPRW*, 2017. 6
- [2] Yuang Ai, Xiaoqiang Zhou, Huaibo Huang, Xiaotian Han, Zhengyu Chen, Quanzeng You, and Hongxia Yang. Dreamclear: High-capacity real-world image restoration with privacy-safe dataset curation. In *NeurIPS*, 2025. 1, 3, 6
- [3] blackforestlabs.ai. Flux, offering state-of-the-art performance image generation, 2024. Accessed: 2024-10-07. 2
- [4] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *CVPR*, 2018. 6
- [5] Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *ICCV*, 2019. 6
- [6] Kelvin CK Chan, Xintao Wang, Xiangyu Xu, Jinwei Gu, and Chen Change Loy. Glean: Generative latent bank for large-factor image super-resolution. In *CVPR*, 2021. 3
- [7] Bin Chen, Gehui Li, Rongyuan Wu, Xindong Zhang, Jie Chen, Jian Zhang, and Lei Zhang. Adversarial diffusion compression for real-world image super-resolution. *arXiv preprint arXiv:2411.13383*, 2024. 3
- [8] 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 *ACM MM*, 2022. 3
- [9] Du Chen, Jie Liang, Xindong Zhang, Ming Liu, Hui Zeng, and Lei Zhang. Human guided ground-truth generation for realistic image super-resolution. In *CVPR*, 2023. 3
- [10] 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 *CVPR*, 2021. 3
- [11] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, et al. Pixart- α : Fast training of diffusion transformer for photorealistic text-to-image synthesis. *arXiv preprint arXiv:2310.00426*, 2023. 2
- [12] Junyang Chen, Jinshan Pan, and Jiangxin Dong. Faithdiff: Unleashing diffusion priors for faithful image super-resolution. In *CVPR*, 2025. 3
- [13] Xiangyu Chen, Xintao Wang, Jiantao Zhou, Yu Qiao, and Chao Dong. Activating more pixels in image super-resolution transformer. In *CVPR*, 2023. 3
- [14] Zheng Chen, Yulun Zhang, Jinjin Gu, Linghe Kong, Xiaokang Yang, and Fisher Yu. Dual aggregation transformer for image super-resolution. In *ICCV*, 2023. 3
- [15] Kun Cheng, Lei Yu, Zhijun Tu, Xiao He, Liyu Chen, Yong Guo, Mingrui Zhu, Nannan Wang, Xinbo Gao, and Jie Hu. Effective diffusion transformer architecture for image super-resolution. In *AAAI*, 2025. 3
- [16] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In *CVPR*, 2023. 4
- [17] Qinpeng Cui, Yixuan Liu, Xinyi Zhang, Qiqi Bao, Zhongdao Wang, Qingmin Liao, Li Wang, Tian Lu, and Emad Barsoum. Taming diffusion prior for image super-resolution with domain shift sdes. *arXiv preprint arXiv:2409.17778*, 2024. 3
- [18] Tao Dai, Jianrui Cai, Yongbing Zhang, Shu-Tao Xia, and Lei Zhang. Second-order attention network for single image super-resolution. In *CVPR*, 2019. 3
- [19] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *ECCV*, 2014. 2, 3
- [20] Linwei Dong, Qingnan Fan, Yihong Guo, Zhonghao Wang, Qi Zhang, Jinwei Chen, Yawei Luo, and Changqing Zou. Tsd-sr: One-step diffusion with target score distillation for real-world image super-resolution. In *CVPR*, 2025. 3
- [21] 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 *ICML*, 2024. 2, 4
- [22] Junhao Gu, Peng-Tao Jiang, Hao Zhang, Mi Zhou, Jinwei Chen, Wenming Yang, and Bo Li. Consissr: Delving deep into consistency in diffusion-based image super-resolution. *arXiv preprint arXiv:2410.13807*, 2024. 3
- [23] Shuhang Gu, Andreas Lugmayr, Martin Danelljan, Manuel Fritsche, Julien Lamour, and Radu Timofte. Div8k: Diverse 8k resolution image dataset. In *ICCVW*, 2019. 6
- [24] Gu Jinjin, Cai Haoming, Chen Haoyu, Ye Xiaoxing, Jimmy S Ren, and Dong Chao. Pipal: a large-scale image quality assessment dataset for perceptual image restoration. In *ECCV*, 2020. 6
- [25] Tero Karras. A style-based generator architecture for generative adversarial networks. *arXiv preprint arXiv:1812.04948*, 2019. 6
- [26] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *ICCV*, 2021. 6
- [27] Dehong Kong, Fan Li, Zhixin Wang, Jiaqi Xu, Renjing Pei, Wenbo Li, and WenQi Ren. Dual prompting image restoration with diffusion transformers. In *CVPR*, 2025. 3
- [28] Hao Li, Xiang Chen, Jiangxin Dong, Jinhui Tang, and Jinshan Pan. Foundir: Unleashing million-scale training data to advance foundation models for image restoration. *arXiv preprint arXiv:2412.01427*, 2024. 3

- [29] Jianze Li, Jiezhong Cao, Zichen Zou, Xiongfei Su, Xin Yuan, Yulun Zhang, Yong Guo, and Xiaokang Yang. Distillation-free one-step diffusion for real-world image super-resolution. *arXiv preprint arXiv:2410.04224*, 2024. 3
- [30] Jianze Li, Jiezhong Cao, Yong Guo, Wenbo Li, and Yulun Zhang. One diffusion step to real-world super-resolution via flow trajectory distillation. *arXiv preprint arXiv:2502.01993*, 2025. 3
- [31] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *ICCV*, 2021. 3, 6
- [32] Jie Liang, Hui Zeng, and Lei Zhang. Details or artifacts: A locally discriminative learning approach to realistic image super-resolution. In *CVPR*, 2022. 3
- [33] Jie Liang, Hui Zeng, and Lei Zhang. Efficient and degradation-adaptive network for real-world image super-resolution. In *ECCV*, 2022. 3
- [34] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPRW*, 2017. 3
- [35] Xinqi Lin, Jingwen He, Ziyang Chen, Zhaoyang Lyu, Bo Dai, Fanghua Yu, Yu Qiao, Wanli Ouyang, and Chao Dong. Diffbir: Toward blind image restoration with generative diffusion prior. In *ECCV*, 2024. 1, 3, 6
- [36] Yong Liu, Hang Dong, Jinshan Pan, Qingji Dong, Kai Chen, Rongxiang Zhang, Lean Fu, and Fei Wang. Patchscaler: An efficient patch-independent diffusion model for super-resolution. *arXiv preprint arXiv:2405.17158*, 2024. 3
- [37] Mehdi Noroozi, Isma Hadji, Brais Martinez, Adrian Bulat, and Georgios Tzimiropoulos. You only need one step: Fast super-resolution with stable diffusion via scale distillation. In *ECCV*, 2024. 3
- [38] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 2
- [39] 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
- [40] Yunpeng Qu, Kun Yuan, Kai Zhao, Qizhi Xie, Jinhua Hao, Ming Sun, and Chao Zhou. Xpsr: Cross-modal priors for diffusion-based image super-resolution. In *ECCV*, 2024. 3
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 4
- [42] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *JMLR*, 2020. 4
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3
- [44] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, 2022. 2
- [45] Haoze Sun, Wenbo Li, Jianzhuang Liu, Haoyu Chen, Renjing Pei, Xueyi Zou, Youliang Yan, and Yujiu Yang. Coser: Bridging image and language for cognitive super-resolution. In *CVPR*, 2024. 3
- [46] Lingchen Sun, Rongyuan Wu, Zhiyuan Ma, Shuaizheng Liu, Qiaosi Yi, and Lei Zhang. Pixel-level and semantic-level adjustable super-resolution: A dual-lora approach. *arXiv preprint arXiv:2412.03017*, 2024. 3
- [47] Lingchen Sun, Rongyuan Wu, Zhengqiang Zhang, Hongwei Yong, and Lei Zhang. Improving the stability of diffusion models for content consistent super-resolution. *CoRR*, 2024. 3
- [48] 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 *CVPRW*, 2017. 6
- [49] Li-Yuan Tsao, Hao-Wei Chen, Hao-Wei Chung, Deqing Sun, Chun-Yi Lee, Kelvin CK Chan, and Ming-Hsuan Yang. Holisdip: Image super-resolution via holistic semantics and diffusion prior. *arXiv preprint arXiv:2411.18662*, 2024. 3
- [50] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, Dhruv Nair, Sayak Paul, William Berman, Yiyi Xu, Steven Liu, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>, 2022. 2, 3, 6
- [51] Yuhao Wan, Peng-Tao Jiang, Qibin Hou, Hao Zhang, Jinwei Chen, Ming-Ming Cheng, and Bo Li. Clearsr: Latent low-resolution image embeddings help diffusion-based real-world super resolution models see clearer. *arXiv preprint arXiv:2410.14279*, 2024. 3
- [52] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *AAAI*, 2023. 6
- [53] Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin CK Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. *IJCV*, 2024. 3, 6
- [54] 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 *ECCVW*, 2018. 3
- [55] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *ICCV*, 2021. 2, 3, 6
- [56] Yufei Wang, Wenhan Yang, Xinyuan Chen, Yaohui Wang, Lanqing Guo, Lap-Pui Chau, Ziwei Liu, Yu Qiao, Alex C Kot, and Bihan Wen. Sinsr: diffusion-based image super-resolution in a single step. In *CVPR*, 2024. 3
- [57] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004. 6
- [58] 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 *ECCV*, 2020. 6
- [59] 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. 3, 6
- [60] Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. Seesr: Towards semantics-aware real-world image super-resolution. In *CVPR*, 2024. 1, 3, 6
- [61] Enze Xie, Junsong Chen, Junyu Chen, Han Cai, Haotian Tang, Yujun Lin, Zhekai Zhang, Muyang Li, Ligeng Zhu, Yao Lu, et al. Sana: Efficient high-resolution image synthesis with linear diffusion transformers. *arXiv preprint arXiv:2410.10629*, 2024. 3
- [62] Liangbin Xie, Xintao Wang, Xiangyu Chen, Gen Li, Ying Shan, Jiantao Zhou, and Chao Dong. Desra: detect and delete the artifacts of gan-based real-world super-resolution models. *arXiv preprint arXiv:2307.02457*, 2023. 3
- [63] Rui Xie, Chen Zhao, Kai Zhang, Zhenyu Zhang, Jun Zhou, Jian Yang, and Ying Tai. Addsr: Accelerating diffusion-based blind super-resolution with adversarial diffusion distillation. *arXiv preprint arXiv:2404.01717*, 2024. 3
- [64] Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assessment. In *CVPR*, 2022. 6
- [65] Tao Yang, Rongyuan Wu, Peiran Ren, Xuansong Xie, and Lei Zhang. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. In *ECCV*, 2024. 3
- [66] 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 *CVPR*, 2024. 1, 6
- [67] Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Resshift: Efficient diffusion model for image super-resolution by residual shifting. In *NeurIPS*, 2024. 3, 6
- [68] Aiping Zhang, Zongsheng Yue, Renjing Pei, Wenqi Ren, and Xiaochun Cao. Degradation-guided one-step image super-resolution with diffusion priors. *arXiv preprint arXiv:2409.17058*, 2024. 3
- [69] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *ICCV*, 2021. 2, 3
- [70] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *CVPR*, 2023. 1, 2, 3
- [71] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 6
- [72] Weixia Zhang, Guangtao Zhai, Ying Wei, Xiaokang Yang, and Kede Ma. Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In *CVPR*, 2023. 6
- [73] Xindong Zhang, Hui Zeng, Shi Guo, and Lei Zhang. Efficient long-range attention network for image super-resolution. In *ECCV*, 2022. 3
- [74] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *ECCV*, 2018.
- [75] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *CVPR*, 2018. 3