

XPSR: Cross-modal Priors for Diffusion-based Image Super-Resolution

Yunpeng Qu^{1,2†}, Kun Yuan^{2†‡}, Kai Zhao²,
Qizhi Xie^{1,2}, Jinhua Hao², Ming Sun² and Chao Zhou²

¹ Tsinghua University, China, Beijing

² Kuaishou Technology, China, Beijing

{qyp21,xqz20}@mails.tsinghua.edu.cn

{yuankun03,zhaokai05,haojinhua,sunming03,zhouchao}@kuaishou.com

Abstract. Diffusion-based methods, endowed with a formidable generative prior, have received increasing attention in Image Super-Resolution (ISR) recently. However, as low-resolution (LR) images often undergo severe degradation, it is challenging for ISR models to perceive the semantic and degradation information, resulting in restoration images with incorrect content or unrealistic artifacts. To address these issues, we propose a *Cross-modal Priors for Super-Resolution (XPSR)* framework. Within XPSR, to acquire precise and comprehensive semantic conditions for the diffusion model, cutting-edge Multimodal Large Language Models (MLLMs) are utilized. To facilitate better fusion of cross-modal priors, a *Semantic-Fusion Attention* is raised. To distill semantic-preserved information instead of undesired degradations, a *Degradation-Free Constraint* is attached between LR and its high-resolution (HR) counterpart. Quantitative and qualitative results show that XPSR is capable of generating high-fidelity and high-realism images across synthetic and real-world datasets. Codes are released at <https://github.com/qyp2000/XPSR>.

Keywords: Image super-resolution · Image restoration · Diffusion models · Multimodal large language models

1 Introduction

The objective of Image Super-Resolution (ISR) entails generating a perceptually authentic high-resolution (HR) image from its low-resolution (LR) counterpart, which is characterized by unknown and intricate degradation processes. Despite the considerable achievements of GAN-based methods [37, 63, 73], they still struggle to reproduce vivid and realistic textures. This is mainly attributable to the domain gap between the synthetic training data and the real-world test data, and the excessive fidelity-oriented optimization objective [67].

Recently, denoising diffusion probabilistic models (DDPMs) are emerging as successors to GANs across a range of generation tasks [23, 47, 52, 76], due

[†] Equal contribution.

[‡] Project leader.

to their strong ability in modeling complicated distributions. Some pioneering works [14, 28] have adopted DDPMs to tackle the ISR problems. Notably, ISR introduces further challenges, on account of high restoration fidelity. This requirement starkly contrasts with the inherent stochasticity of DDPMs.

Given that the large-scale pretrained text-to-image (T2I) models (*e.g.* StableDiffusion (SD) [52]), trained on datasets surpassing 5 billion image-text pairs and embodying robust and plentiful natural image priors, other methods [39, 61, 67, 69] utilize ControlNet [76] to harness pre-trained SD priors for ISR. Encouragingly, these methods have exhibited the incredible capability to generate realistic image details. StableSR [61] and DiffBIR [39] straightforwardly set the prompt condition to empty, relying on extracting semantic information from LR images. However, LR images undergo complex degradations, making it challenging for ISR models to extract semantic priors, which are provided in the form of textual prompts within pre-trained T2I models. PASD [69] and SeeSR [67] leverage off-the-shelf tagging models for the extraction of object labels as high-level prompts. However, these tagging prompts lack more complex information such as spatial location and scene understanding, which are crucial for generating comprehensive images in T2I models. These prompts also fail to capture the inherent distortions within images, yet such low-level priors facilitate the modeling of the degradation process, thereby enabling clearer restorations in ISR [3, 74].

In this paper, we explore the significance of different semantic priors from multi-modal large language models (MLLMs) for ISR. Based on the analysis, we present a *Cross-modal Priors for Super-Resolution (XPSR)* framework, utilizing cross-modal priors to guide diffusion models in generating more high-fidelity and realistic images. To furnish more precise and perceptually aligned semantic priors, we introduce MLLMs to offer multi-level prompts of LR images. To effectively fuse cross-model semantic priors with generative priors, we propose a *Semantic-Fusion Attention (SFA)* for adaptive selection of semantic features. To enhance the extraction of semantic-preserved features, we further devise a *Degradation-Free Constraint (DFC)*. Our **contributions** are as follows:

1. We explore the significance of semantic priors in diffusion-based ISR, wherein high-level priors offer a wealth of semantic information while low-level priors assist in modeling degradation mechanisms. Furthermore, cutting-edge MLLMs are utilized to obtain the desired appropriate priors.
2. We propose the XPSR, where the SFA is adopted to fuse multi-level semantic priors with the diffusion model in a parallel cross-attention manner.
3. To obtain semantic-preserved features instead of degradations, the DFC is attached between the LR and HR images in the pixel and latent space.
4. Through quantitative and qualitative analysis, XPSR demonstrates a strong capability in generating high-fidelity and high-realism images, achieving **the best performance** on multiple image quality metrics across several datasets. Extensive ablation studies prove the validity of each component.

2 Related Work

2.1 Degradation Prior-based Image Super-resolution

Initial investigations [8, 10, 12, 13, 38, 81] into ISR mainly focused on the restoration of LR images through pre-determined degradation types, such as bicubic downsampling, blurring, noise, and other factors. Yet, their efficacy markedly diminishes in practical settings due to their limited adaptability. [17, 45, 79]. To elevate performance further, recent advancements (*i.e.*, GAN-based methods [7, 36, 37, 63, 73]) have delved into more complex degradation models to mimic the real-world distortions closely. Among them, BSRGAN [73] proposed using random shuffling combinations of basic degradation operations. Real-ESRGAN [63] employed a high-order process to simulate degradation in real-life procedures, *e.g.*, cameras, image editing, transmission, *etc.* Beyond the dependence on explicit priors, other studies [6, 18, 35, 43, 44] also discovered adopting implicit priors through adversarial training. However, these methods fall short, chiefly due to insufficient prior knowledge, resulting in suboptimal output quality.

2.2 Diffusion Prior-based Image Super-resolution

Analogous to GAN-based methodologies, diffusion models [4, 24, 47, 51, 53, 54, 76] have earned escalating recognition for their superior ability to embody generative priors. Of late, these models have been skillfully utilized as generative priors within the realm of ISR as well [29, 55, 56, 62]. Among them, StableSR [61] trained a time-aware encoder to fine-tune the StableDiffusion (SD) model [52] and employed feature warping to balance between fidelity and perceptual quality. DiffBIR [39] utilized a two-stage approach, initially reconstructing the image as a preliminary estimate, and subsequently leveraging the SD to refine and augment visual details. To effectively utilize the potential of pretrained text-to-image diffusion models, PASD [69] and SeeSR [67] introduced off-the-shelf high-level semantic information (*i.e.*, object tags) as extra conditions. Through these endeavors, ISR can reproduce more realistic image details. Nevertheless, when confronted with LR images that exhibit complicated distortions, and intertwined spatial positions among objects, the generated result will be substandard.

2.3 Cross-modal Semantic Prior-based Image Super-resolution

Textual prompts are essential in guiding targeted image generation in T2I diffusion models [5, 21, 30]. Previous ISR works [67, 69] have attempted to utilize existing visual perception models (*i.e.*, ResNet [20], BLIP [34], RAM [80]) to understand image content and provide cross-modal semantic priors as prompts. However, these prompts fall short as they only encompass basic object recognition, lacking high-level information (*e.g.*, spatial positioning and scene understanding) and low-level information (*e.g.*, image quality and sharpness). Leveraging CLIP [50] and additional adapters to align visual inputs with textual inputs into large-scale language models (LLMs, *e.g.*, GPT-4 [1], LLaMA [59]),

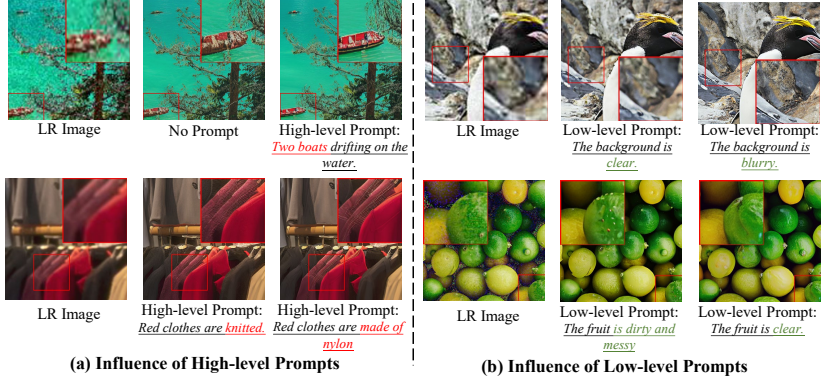


Fig. 1: Impact of textual prompts for conditional T2I diffusion models. **(a)** When given high-level prompts containing object categories and detailed textures, high-realism HR images can be restored based on ambiguous LR images. **(b)** When furnished with **accurate** low-level prompts that encompass distortion types or the general quality of LR images, high-fidelity images can be generated from blurry or noisy inputs.

MLLMs [16, 33, 40, 49, 77, 83] exhibit remarkable visual comprehension capabilities. By scaling up parameters and training data, MLLMs have shown robust performance in both high-level and low-level perceptual tasks, enabling the extraction of cross-model semantic priors that align with human perception [66, 70, 71].

3 Methodology

3.1 Cross-modal Semantic Priors

Current diffusion-based ISR methods typically capture semantic structures from LR images and utilize the T2I model to generate realistic HR images. However, LR images often undergo complex degradation processes, complicating the task for ISR models to extract semantic content independently. This difficulty is exacerbated by constraints in scale and the availability of training data, leading to the restoration of blurry and indistinguishable images. Therefore, it is necessary to provide additional conditions to the diffusion model more effectively. Considering that T2I models fundamentally generate images guided by textual input, we believe that leveraging textual prompts as semantic priors is a viable approach to enhance the capabilities of T2I models for ISR.

Based on common cognition [15], the semantic information contained in images can be roughly divided into two categories: *high-level* and *low-level semantics*. To more intuitively reflect the impact of different types of prompts on diffusion-based ISR tasks, we have pre-visualized some results in Fig. 1. **First**, we provide *high-level semantic priors* related to the LR images, which may encompass information about the objects in the image, spatial positioning, scene descriptions, and more. As illustrated in Fig. 1 (a), it becomes possible to restore images that are clearer and retain the corresponding semantic information

(e.g., the boat and the clothes). These show the advantage of utilizing additional high-level semantic priors upon relying on sole diffusion models. **Second**, we use *low-level semantic priors* as guidance, which includes the perception of overall quality, sharpness, noise level, and other distortions about the LR image. As shown in Fig. 1 (b), accurate low-level conditions are necessary for generating high-quality images. For the 1st case, where the background of the LR image is highly blurred, providing the correct prompt (i.e., "The background is blurry") enables the restoration of a clear image. Conversely, an incorrect prompt (i.e., "The background is clear") may lead to the restoration of blurry details due to the model's assumption of sufficient clarity in the LR image. For the 2nd case, an incorrect prompt (i.e., "The fruit is messy and blurry") leads to unrealistic artifacts. This implies that low-level semantic priors assist in modeling the degradation process, which benefits clearer images. Therefore, we believe that by combining both high- and low-level semantic priors, restoration results can be semantically accurate and rich in detail as well.

3.2 Framework of XPSR

To combine semantic priors with the LR image, as conditions of the diffusion model, there still exist three problems that require immediate resolution: (1) *how to obtain semantic priors that describe accurately?* (2) *How to effectively incorporate these priors with well-defined diffusion models?* (3) *How to extract semantic-preserved but degradation-unrelated conditions from LR images?*

To address these problems, we propose a framework called *Cross-modal Priors for Super Resolution (XPSR)*. As depicted in Fig. 2, it consists of two main stages: generation of semantic priors and image restoration with these priors. In the first stage (Fig. 2 (a)), **to address the first problem**, we rely on existing SOTA MLLMs to acquire high-level and low-level semantic priors for the LR image (Sec. 3.3). Then two varieties of embeddings are derived upon input into the CLIP [50] text encoder. In the second stage (Fig. 2 (b)), we employ the powerful pretrained T2I SD model [52] as the backbone. Additionally, we utilize ControlNet [76] as a controller, guiding the conditional image restoration process. We clone the encoder of the Unet architecture from SD as a trainable copy to initialize ControlNet. **To address the second problem**, we design a *Semantic-Fusion Attention (SFA)* module, which facilitates the interaction between the semantic priors obtained in the first stage and the generated priors from the T2I model (Sec. 3.4). **To address the third problem**, during training, we further attach a *Degradation-Free Constraint (DFC)* to ControlNet, reducing the impact of degradation and extracting semantic-aware representations from the LR image (Sec. 3.5). During inference, only ControlNet and Unet parts, along with the MLLM, are used for predicting restoration images.

3.3 Semantic Prompts from MLLM

By scaling up the size of data and the scale of models, MLLMs have demonstrated impressive capabilities in semantic comprehension [70]. Accordingly, we employ

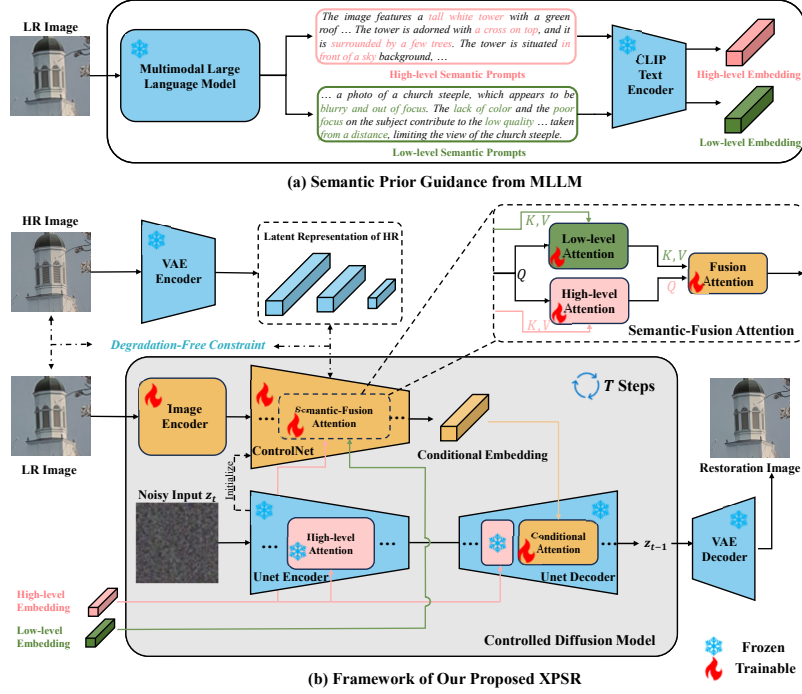


Fig. 2: Framework of XPSR. **(a)** First, we integrate an MLLM to acquire semantic priors, encompassing both high-level and low-level descriptions for the LR image. Two varieties of embeddings are derived upon input into the CLIP text encoder. **(b)** Next, the LR image, along with the embeddings above, are input into the controlled diffusion model as conditions through a *Semantic-Fusion Attention (SFA)*, adhering to the defined workflow. Besides, a *Degradation-Free Constraint (DFC)* is appended to the ControlNet part, alleviating the challenge of discerning distortions.

a noteworthy MLLM called LLaVA [40] to perceive the LR images and extract semantic priors. By defining appropriate instructions, we aim to facilitate its understanding of semantics from both high- and low-level perspectives.

Determining appropriate instructions to guide the MLLMs in generating outputs that align with human preferences is a matter that requires iterative practice [9]. *In the acquisition of high-level semantic priors*, we have ultimately opted for the instruction "Please provide a descriptive summary of the content of this image" as the guidance for LLaVA. In our practical experience, we observe that similar instructions consistently yield similar results with LLaVA, generating the intended high-level semantics, including object descriptions, spatial locations, scenes, and other relevant content. *In the acquisition of low-level semantic priors*, we first analyze the distribution of degradations in real-world scenarios. The degradation primarily occurs due to processes such as capturing, transcoding compression, or transmission. Consequently, it introduces

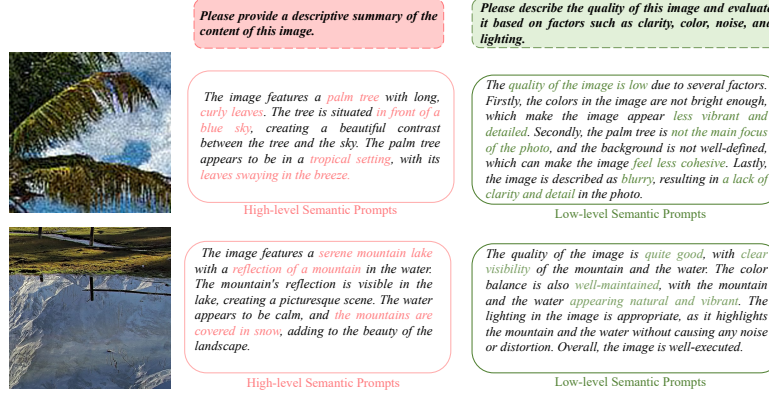


Fig. 3: Given appropriate instructions, LLaVA can generate high- and low-level semantic prompts consistent with human perception for both high- and low-quality images.

various distortions such as motion blur, low-light conditions, color shifts, noise, and compression artifacts [26, 26, 41, 46, 57, 72, 75, 82]. Therefore, by incorporating guidance in the instructions on these quality-related factors, we observe that this approach generates more realistic and detailed descriptions. We adopt the instruction "Please describe the quality of this image and evaluate it based on factors such as clarity, color, noise, and lighting" to generate descriptions related to overall quality, clarity, noise level, color, and other relevant information. As shown in Fig. 3, LLaVA is capable of providing high-level and low-level semantic descriptions that align with human perception.

3.4 Semantic Priors Fusion for Diffusion Model

To integrate extra semantic priors into the diffusion model, an intuitive and naive approach is to employ two consecutive cross-attention (*i.e.*, attention type for the original T2I) structures as a serial connection into the ControlNet. Given the feature \mathbf{x}_k in the k -th layer, the computation process can be noted as:

$$\mathbf{x}_{k+1} = \mathcal{CA}_l(\mathcal{CA}_h(\mathbf{x}_k, c_h), c_l), \quad (1)$$

where \mathcal{CA}_h and \mathcal{CA}_l represent the high-level and low-level cross-attention while c_h and c_l denote the two types of semantic prompts, respectively. However, as these two types of priors are distinct, serial processing may cause a certain part of the information to be overwritten and obtain suboptimal results.

Therefore, we design a new *Semantic-Fusion Attention (SFA)*, which employs two parallel branches of cross-attention, and then fuses them to obtain the final condition using a fusion attention \mathcal{CA}_f . As shown in Fig. 2, the fusion attention obtains the query Q from the high-level results, and the key K and value V from the low-level results. The parallel type of SFA can be written as:

$$\mathbf{x}_{k+1} = \mathcal{CA}_f(\mathcal{CA}_h(\mathbf{x}_k, c_h), \mathcal{CA}_l(\mathbf{x}_k, c_l)). \quad (2)$$

In this way, SFA can achieve a balance between priors from different levels, allowing for an adaptive selection of semantic features. Ablation in Tab. 4 further verifies the effectiveness of this setting. Specifically, we claim that a low-level understanding of LR images is not necessary for Unet, which receives noise as input, so only high-level attention is used. To enhance the generalizability of SD for downstream tasks and fortify the fusion of features derived from LR images, we add a cross-attention module termed *Conditional Attention* in the Unet.

3.5 Degradation-Free Constraint

Real-world images may experience a variety of degradations, such as blurriness, blocky artifacts, *etc.*, leading to distortions that affect both high-frequency and low-frequency components (appeared in the pixel- and latent-space) [26, 57]. Aiming to mitigate the impact of degradations and extract robust semantic information from LR images, we propose a *Degradation-Free Constraint (DFC)* which attaches constraints on both *the pixel space* and *the latent space*.

As shown in Fig. 2, the LR image is initially passed through a pyramid-shaped image encoder in ControlNet for downsampling. Following [69], we apply a pixel-space constraint in the image encoder. At the i -th layer of the pyramid encoder, we employ a single-layer convolution to map the feature map into an image $\hat{x}_i \in \mathbb{R}^{H_i \times W_i \times 3}$ with RGB channels. As the scale is reduced by half compared to the previous layer, we apply an L_1 loss to make \hat{x}_i as close as possible to the downsampled result of the original HR image $x_{hr,i}$, denoted as $\|x_{hr,i} - \hat{x}_i\|_1$.

Besides, ControlNet adopts a pyramid-like Unet Encoder to encode semantic features at the latent level, we further apply a similar latent-space constraint. The feature map obtained from the j -th layer is mapped to the latent space as \hat{z}_j . We align it with the downsampled result of the HR latent $z_{hr,j}$ at scale j using the L_1 loss $\|z_{hr,j} - \hat{z}_j\|_1$. The final DFC is a combination of constraints in both the pixel space and the latent space, which can be noted as:

$$\mathcal{L}_{DFC} = \sum_{i=1}^3 \|x_{hr,i} - \hat{x}_i\|_1 + \sum_{j=1}^3 \|z_{hr,j} - \hat{z}_j\|_1. \quad (3)$$

3.6 Training and Testing Strategy

During training, the HR image is mapped to the latent embedding z_{hr} . We add noise ϵ through the diffusion process by following t steps to obtain the noisy latent z_{hr}^t , where t is randomly sampled from a range of $[1, T]$. Overall, XPSR relies on the LR image x_{lr} , noisy latent z_{hr}^t , high-level prompt c_h , and low-level prompt c_l to predict the added noise. The optimization objective is:

$$\mathcal{L}_D = \mathbb{E}_{z_{hr}, t, c_h, c_l, x_{lr}, \epsilon \sim \mathcal{N}(0,1)} \left[\left\| \epsilon - \epsilon_\theta(z_{hr}^t, t, c_h, c_l, x_{lr}) \right\|_2^2 \right]. \quad (4)$$

where $\epsilon_\theta(\cdot)$ represents the mapping function of XPSR. The final loss function is a weighted sum of the diffusion loss and the DFC loss:

$$\mathcal{L} = \mathcal{L}_D + \lambda \mathcal{L}_{DFC}, \quad (5)$$

where λ is a balancing coefficient. To reduce training costs and leverage the generative prior of SD, we freeze all the parameters of SD during training and only train the ControlNet and the added *Conditional Attention*.

During testing, we employ the classifier-free guidance strategy [25], which allows the diffusion model to generate better-quality images through negative prompts without additional training. In each step, we rely on regular high-level prompts c_h for prediction while simultaneously replacing c_h with negative prompts c_{neg} . These predictions are fused to obtain the final output:

$$\begin{aligned}\hat{e} &= \epsilon_{\theta}(z_{lr}^t, t, c_h, c_l, x_{lr}) \\ \hat{e}_{neg} &= \epsilon_{\theta}(z_{lr}^t, t, c_{neg}, c_l, x_{lr}) \\ \tilde{e} &= \hat{e} + \lambda_s(\hat{e} - \hat{e}_{neg})\end{aligned}\tag{6}$$

where λ_s is the guidance scale and z_{lr}^t is the noisy latent of LR image. In practice, We employ combinations of negative words "*blurry, dotted, noise, unclear, low-res, over-smoothed*" as negative prompts to generate higher-quality images.

4 Experiments

4.1 Experimental Settings

Training and testing datasets. We train our XPSR on DIV2K [2], DIV8K [19], Flickr2K [58], Unsplash2K [32], OST [64], and the first 5K face images from FFHQ [27]. We use the degradation pipeline of Real-ESRGAN [63] to synthesize LR-HR image pairs for training. To conduct a comprehensive evaluation of ISR, we conduct testing using synthetic and real-world datasets. The synthetic dataset is generated from the DIV2K validation set, where we randomly crop 3K patches following the same degradation pipeline. Following [61, 67], the LR image are center-cropped to 128×128 . For the real-world dataset, we employ the DrealSR and RealSR datasets for evaluation, with each image being center-cropped. Each HR image in the training and testing sets has a resolution of 512×512 .

Implementation details. We employ the pretrained SD-v1.5³ as the base T2I model. For training, we utilize an AdamW [42] optimizer with a weight decay of $1e-2$ to finetune XPSR for 100K iterations. The batch size and the learning rate are set to 32 and $5e-5$. All experiments are conducted on 8 NVIDIA A800 GPUs. For inference, we adopt DDPM sampling [48] with 20 timesteps. The balancing coefficient λ in Eq. (5) is 0.05 and the guidance scale λ_s in Eq. (6) is 5.5.

Evaluation metrics. We employ a range of widely used reference and non-reference metrics to conduct a comprehensive quantitative evaluation of ISR methods. In reference-based metrics, PSNR and SSIM [65] (calculated on the Y channel in YCbCr space) are fidelity metrics, while LPIPS [78], DISTS [11] are quality evaluation metrics. FID [22] calculates the distance between the distributions of generated images and reference images. MANIQA [68], MUSIQ [31], and CLIPSIQA [60] are non-reference image quality assessment (IQA) metrics.

³ <https://huggingface.co/runwayml/stable-diffusion-v1-5>




Fig. 4: Drawbacks of current full-reference metrics (*e.g.*, PSNR, SSIM, LPIPS). The visualization results show that XPSR generates images with high-realism and high-fidelity in human perception, but obtains lower scores in some cases.

Table 1: Quantitative comparison with SOTA methods on both synthetic and real-world benchmarks. Red and blue colors are the best and second-best performance.

| Dataset | Metrics | GAN-based SR | | | | Diffusion-based SR | | | | |
|-----------|--------------------|--------------|-------------|--------|--------|--------------------|---------|--------|--------|--------|
| | | BSRGAN | Real-ESRGAN | SwinIR | LDM | StableSR | DiffBIR | PASD | SeeSR | XPSR |
| DIV2K-Val | PSNR \uparrow | 24.42 | 24.30 | 23.77 | 21.66 | 23.26 | 23.49 | 23.59 | 23.56 | 22.80 |
| | SSIM \uparrow | 0.6164 | 0.6324 | 0.6186 | 0.4752 | 0.5670 | 0.5568 | 0.5899 | 0.5981 | 0.5627 |
| | LPIPS \downarrow | 0.3511 | 0.3267 | 0.3910 | 0.4887 | 0.3228 | 0.3638 | 0.3611 | 0.3283 | 0.3761 |
| | DISTS \downarrow | 0.2369 | 0.2245 | 0.2291 | 0.2693 | 0.2116 | 0.2177 | 0.2134 | 0.2008 | 0.2217 |
| | FID \downarrow | 50.99 | 44.34 | 44.45 | 55.04 | 28.32 | 34.55 | 39.74 | 28.89 | 33.38 |
| | MANIQA \uparrow | 0.3547 | 0.3756 | 0.3411 | 0.3589 | 0.4173 | 0.4598 | 0.4440 | 0.5046 | 0.6080 |
| | CLIPQA \uparrow | 0.5253 | 0.5205 | 0.5213 | 0.5570 | 0.6752 | 0.6731 | 0.6573 | 0.6959 | 0.7816 |
| | MUSIQ \uparrow | 60.18 | 59.76 | 57.21 | 57.46 | 65.19 | 65.57 | 66.58 | 68.35 | 69.99 |
| RealSR | PSNR \uparrow | 26.38 | 25.68 | 25.88 | 25.66 | 24.69 | 24.94 | 25.21 | 25.31 | 24.19 |
| | SSIM \uparrow | 0.7651 | 0.7614 | 0.7671 | 0.6934 | 0.7090 | 0.6664 | 0.7140 | 0.7284 | 0.6870 |
| | LPIPS \downarrow | 0.2656 | 0.2710 | 0.2614 | 0.3367 | 0.3003 | 0.3485 | 0.2986 | 0.2993 | 0.3517 |
| | DISTS \downarrow | 0.2124 | 0.2060 | 0.2061 | 0.2324 | 0.2134 | 0.2257 | 0.2125 | 0.2224 | 0.2471 |
| | FID \downarrow | 141.25 | 135.14 | 132.80 | 133.34 | 131.72 | 127.59 | 139.42 | 126.21 | 141.95 |
| | MANIQA \uparrow | 0.3763 | 0.3736 | 0.3561 | 0.3375 | 0.4167 | 0.4378 | 0.4418 | 0.5370 | 0.6059 |
| | CLIPQA \uparrow | 0.5114 | 0.4487 | 0.4433 | 0.6053 | 0.6200 | 0.6396 | 0.6009 | 0.6638 | 0.7354 |
| | MUSIQ \uparrow | 63.28 | 60.37 | 59.28 | 56.32 | 65.25 | 64.32 | 66.61 | 69.56 | 70.23 |
| DRealSR | PSNR \uparrow | 28.70 | 28.61 | 28.20 | 27.78 | 27.87 | 26.57 | 27.45 | 28.13 | 26.62 |
| | SSIM \uparrow | 0.8028 | 0.8052 | 0.7983 | 0.7152 | 0.7427 | 0.6516 | 0.7539 | 0.7711 | 0.7220 |
| | LPIPS \downarrow | 0.2858 | 0.2819 | 0.2830 | 0.3745 | 0.3333 | 0.4537 | 0.3331 | 0.3142 | 0.3864 |
| | DISTS \downarrow | 0.2144 | 0.2089 | 0.2103 | 0.2417 | 0.2297 | 0.2724 | 0.2322 | 0.2230 | 0.2606 |
| | FID \downarrow | 155.62 | 147.66 | 146.38 | 164.87 | 148.18 | 160.67 | 173.40 | 147.00 | 164.68 |
| | MANIQA \uparrow | 0.3441 | 0.3435 | 0.3311 | 0.3342 | 0.3897 | 0.4602 | 0.4551 | 0.5077 | 0.5713 |
| | CLIPQA \uparrow | 0.5061 | 0.4525 | 0.4522 | 0.5984 | 0.6321 | 0.6445 | 0.6365 | 0.6893 | 0.7360 |
| | MUSIQ \uparrow | 57.16 | 54.27 | 53.01 | 51.37 | 58.72 | 61.06 | 63.69 | 64.75 | 67.84 |

4.2 Comparisons with SOTA Methods

To verify the effectiveness, we compare XPSR with other SOTA GAN-based and Diffusion-based ISR methods⁴, *i.e.*, BSRGAN [73], Real-ESRGAN [63], SwinIR-GAN [37], LDM [4], StableSR [61], DiffBIR [39], PASD [69] and SeeSR [67].

Quantitative comparisons. In Tab. 1, some observations can be found. **First**, XPSR outperforms other SOTA methods in MAINIQA, CLIPQA, and MUSIQ across all datasets by large margins, reflecting excellent restoration results. For example, XPSR surpasses the second-best method SeeSR by **10.34%**, **8.57%**, and **1.64%** in DIV2K-val, respectively. **Second**, Diffusion-based methods generally fall behind GAN-based methods in reference metrics (*e.g.*, PSNR, SSIM).

⁴ All methods are tested based on their official code and models.




Fig. 5: Qualitative comparisons with different SOTA methods. XPSR is adept at precisely restoring the textures and details of specific objects, such as hair and buildings, even under challenging conditions of degradation. **Zoom in for a better view.**

This is mainly because Diffusion-based methods can produce more realistic details but at the expense of fidelity to the LR images. Furthermore, we give several examples that highlight the limitations of current reference metrics in Fig. 4. Our restoration results exhibit **higher quality in human perception, yet they lag in reference metrics**. This phenomenon has also been confirmed in many previous studies [61, 67, 69]. We believe there is a need to update the metrics used to evaluate ISR, aiming to closely mimic human perception.

Qualitative comparisons. In Fig. 5, we present visual comparisons on the test set, revealing key observations. **First**, diffusion-based methods can generate more realistic results than GAN-based methods, showing the advantage of inherent *generative priors*. **Second**, XPSR surpasses StableSR in generating results that are semantically accurate and rich in detail (*e.g.*, the 1st, 2nd, 4th, and 5th row), exhibiting the benefit of employing high-level semantic conditions. **Third**, XPSR also outperforms PASD and SeeSR in detail generation and degradation removal. In the 2nd row, XPSR distinguishes itself as **the sole method** adept at understanding the semantic context of the LR image, thus producing detailed and lucid depictions of the house. Moreover, in the 1st, 4th, and 5th rows, XPSR can

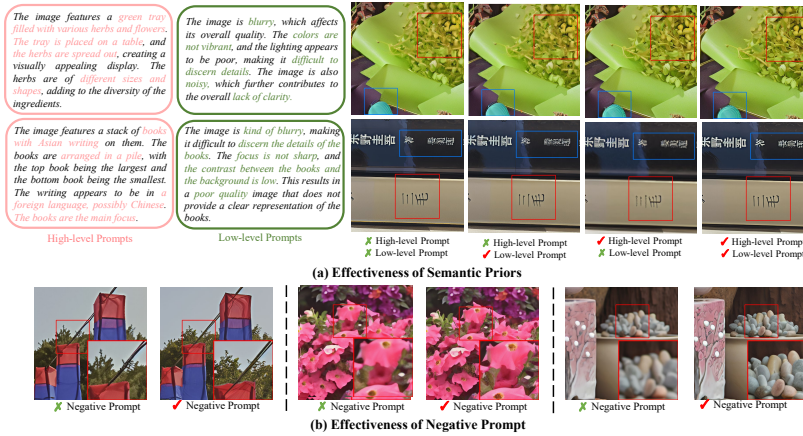


Fig. 6: Effectiveness of cross-modal semantic priors and negative prompts.

Table 2: Ablation on the types of semantic priors.

| Semantic Priors | | <i>DrealSR</i> | | | | | <i>RealSR</i> | | | | |
|-----------------|-----------|----------------|--------|--------|---------|--------|---------------|--------|--------|---------|--------|
| high-level | low-level | SSIM↑ | LPIPS↓ | FID↓ | MANIQA↑ | MUSIQ↑ | SSIM↑ | LPIPS↓ | FID↓ | MANIQA↑ | MUSIQ↑ |
| ✗ | ✗ | 0.7157 | 0.4033 | 188.78 | 0.6078 | 67.31 | 0.6767 | 0.3739 | 157.73 | 0.6426 | 70.95 |
| ✗ | ✓ | 0.7114 | 0.4145 | 195.37 | 0.6213 | 68.73 | 0.6710 | 0.3796 | 162.68 | 0.6438 | 71.05 |
| ✓ | ✗ | 0.7254 | 0.3776 | 160.78 | 0.5505 | 65.79 | 0.6938 | 0.3461 | 139.54 | 0.5752 | 68.74 |
| ✓ | ✓ | 0.7220 | 0.3864 | 164.68 | 0.5713 | 67.84 | 0.6870 | 0.3517 | 141.95 | 0.6059 | 70.23 |

produce more detailed animal hair, clearer leaves, and walnut textures, respectively. These examples highlight the benefits of utilizing *cross-modal semantic priors* from MLLM over-relying on dispersed object tagging information. We are equally impressed by the generative capabilities of all diffusion-based methods, as the outcomes in row 6 show their tremendous potential in the realm of ISR.

4.3 Ablation Study

Effectiveness of semantic priors. To assess the influence of cross-modal semantic priors, the trained XPSR is utilized for inference across various input conditions. The results are given in Tab. 2 and Fig. 6. **First**, the absence of high-level priors leads to a substantial decrease in fidelity metrics (*i.e.*, SSIM, FID). This can be explained that the loss of high-level semantics promotes SFA to overly rely on low-level priors for degradation removal. Although this may yield some gains to MANIQA, it will lead to discrepancies in the content of the restoration images as given in Fig. 6 (a). **Second**, the absence of low-level priors significantly worsens the quality metrics (*i.e.*, MANIQA, MUSIQ). This occurs as SFA gives precedence to high-level priors, leading to a marginal enhancement in SSIM, yet it may produce images that are excessively blurry or marred by noise and artifacts. **Overall**, the integration of dual-level priors facilitates balancing, permitting the restorations that preserve semantic details while also showcasing vivid content.

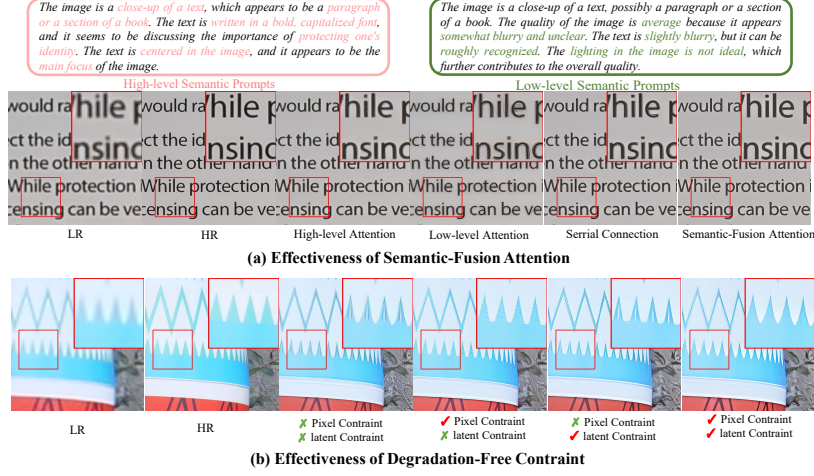


Fig. 7: Effectiveness of *Semantic-Fusion Attention* and *Degradation-Free Constraint*.

Table 3: Ablation on the SFA using different fusion types.

| Semantic-Fusion Methods | <i>DrealSR</i> | | | | | <i>RealSR</i> | | | | |
|-------------------------|----------------|---------------|---------------|---------------|--------------|---------------|---------------|---------------|---------------|--------------|
| | SSIM↑ | LPIPS↓ | FID↓ | MANIQA↑ | MUSIQ↑ | SSIM↑ | LPIPS↓ | FID↓ | MANIQA↑ | MUSIQ↑ |
| High-level | 0.7298 | 0.3674 | 160.04 | 0.5472 | 65.04 | 0.6878 | 0.3500 | 134.68 | 0.5911 | 68.89 |
| Low-level | 0.7578 | 0.3616 | 161.04 | 0.5034 | 63.97 | 0.7047 | 0.3465 | 144.71 | 0.5305 | 66.87 |
| Serial connection | 0.6834 | 0.4134 | 155.23 | 0.5614 | 65.99 | 0.6526 | 0.3908 | 137.49 | 0.6034 | 69.69 |
| SFA(Ours) | 0.7220 | 0.3864 | 164.68 | 0.5713 | 67.84 | 0.6870 | 0.3517 | 141.95 | 0.6059 | 70.23 |

Effectiveness of Semantic-Fusion Attention. Besides the adopted parallel cross-attention mechanisms in SFA, we further test using: (1) only high-level cross-attention, (2) only low-level cross-attention, and (3) a serial concatenation cross-attention. As shown in Tab. 3 and Fig. 7 (a), the parallel version achieves the highest results (*i.e.*, generating clear text without any artifacts). **First**, it further verifies the effectiveness of jointly utilizing high-level and low-level semantic priors to enhance ISR. **Second**, this confirms the efficacy of employing multiple types of conditions in parallel as opposed to a straightforward serial approach.

Effectiveness of Degradation-Free Constraint. We verify different forms of constraints. As shown in Tab. 4, both pixel-level and latent-level constraints contribute to generating more realistic images, resulting in improvements across multiple metrics. We claim that these two constraints aid in preventing the generation of confusing and incomprehensible content. As depicted in Fig. 7 (b), the application of DFC significantly reduces image noise and alleviates the presence of aliasing artifacts. The output images exhibit textures with smooth edges, highlighting the effectiveness of DFC.

Effectiveness of negative prompts. We validate the impact of negative prompts, as described in Sec. 3.6, by setting negative prompts as null text. In Tab. 5a, it

Table 4: Ablation on the *Degradation-Free Constraint*.

| Degradation-Free Constraint | | <i>DrealSR</i> | | | | | <i>RealSR</i> | | | | |
|-----------------------------|--------|----------------|--------|--------|---------|--------|---------------|--------|--------|---------|--------|
| Pixel | Latent | SSIM↑ | LPIPS↓ | DISTS↓ | MANIQA↑ | MUSIQ↑ | SSIM↑ | LPIPS↓ | DISTS↓ | MANIQA↑ | MUSIQ↑ |
| ✗ | ✗ | 0.6979 | 0.3943 | 0.2649 | 0.5684 | 67.50 | 0.6341 | 0.3677 | 0.2525 | 0.6043 | 70.28 |
| ✓ | ✗ | 0.7095 | 0.3938 | 0.2610 | 0.5445 | 64.29 | 0.6783 | 0.3629 | 0.2610 | 0.6042 | 69.91 |
| ✗ | ✓ | 0.7066 | 0.3848 | 0.2580 | 0.5774 | 67.62 | 0.6681 | 0.3665 | 0.2537 | 0.6000 | 70.08 |
| ✓ | ✓ | 0.7220 | 0.3864 | 0.2606 | 0.5713 | 67.84 | 0.6870 | 0.3517 | 0.2471 | 0.6059 | 70.23 |

Table 5: Ablation on the negative prompts and the base T2I diffusion models.

| Negative Prompts | <i>DrealSR</i> | | | | | <i>RealSR</i> | | | | |
|------------------|----------------|--------|--------|---------|--------|---------------|--------|--------|---------|--------|
| | SSIM↑ | LPIPS↓ | FID↓ | MANIQA↑ | MUSIQ↑ | SSIM↑ | LPIPS↓ | FID↓ | MANIQA↑ | MUSIQ↑ |
| ✗ | 0.7355 | 0.3731 | 157.32 | 0.4778 | 61.70 | 0.6987 | 0.3419 | 134.22 | 0.5241 | 65.89 |
| ✓ | 0.7220 | 0.3864 | 164.68 | 0.5713 | 67.84 | 0.6870 | 0.3517 | 141.95 | 0.6059 | 70.23 |

(a) Ablation on the negative prompts.

| Pretrained Model | <i>DrealSR</i> | | | | | <i>RealSR</i> | | | | |
|----------------------|----------------|--------|--------|---------|--------|---------------|--------|--------|---------|--------|
| | SSIM↑ | LPIPS↓ | FID↓ | MANIQA↑ | MUSIQ↑ | SSIM↑ | LPIPS↓ | FID↓ | MANIQA↑ | MUSIQ↑ |
| stable-diffusion-2-1 | 0.7378 | 0.3632 | 155.09 | 0.4910 | 63.79 | 0.6912 | 0.3329 | 138.12 | 0.5148 | 67.10 |
| stable-diffusion-1.5 | 0.7220 | 0.3864 | 164.68 | 0.5713 | 67.84 | 0.6870 | 0.3517 | 141.95 | 0.6059 | 70.23 |

(b) Ablation on the base T2I diffusion models.

significantly improves the perceptual quality based on MAINIQA and MUSIQ. We also provide some visual results in Fig. 6 (b). By applying negative prompts, the images exhibit clearer textures, resulting in more realistic and higher quality.

Different base diffusion models. To explore the impact of the generative prior in the pre-trained T2I model, we analyze different versions of StableDiffusion⁵ in Tab. 5b. Observations indicate a performance drop when employing SD-v2.1. We speculate that this is due to a mismatch in the pre-training resolution, where SD-v2.1 is 768×768 (different from our setting) and SD-v1.5 is 512×512 . However, this does not imply that the pre-training model has achieved its saturation point. **XPSR is orthogonal to these foundational models.** As the industry introduces superior base models, further improvements can be explored.

5 Conclusion

To address the challenge of accurately restoring semantic details in ISR tasks, we propose the XPSR framework. We explore the role of different semantic priors in ISR and propose leveraging MLLMs to extract accurate cross-modal semantic priors as textual prompts. To facilitate the integration of semantic priors, the SFA module is proposed. To retain semantic-preserved information from degraded images, the DFC is applied. Extensive experiments validate the strong performance of XPSR in generating images with clear details and semantic fidelity. We hope this work can inspire exploring the cross-model priors in ISR.

⁵ <https://huggingface.co/stabilityai/stable-diffusion-2-1>

References

1. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
2. Agustsson, E., Timofte, R.: NTIRE 2017 challenge on single image super-resolution: Dataset and study. In: CVPR Workshops. pp. 1122–1131. IEEE Computer Society (2017)
3. Bell-Kligler, S., Shocher, A., Irani, M.: Blind super-resolution kernel estimation using an internal-gan. *Advances in Neural Information Processing Systems* **32** (2019)
4. Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10674–10685. IEEE (2022)
5. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: CVPR. pp. 18392–18402. IEEE (2023)
6. Chan, K.C.K., Wang, X., Xu, X., Gu, J., Loy, C.C.: GLEAN: generative latent bank for large-factor image super-resolution. In: CVPR. pp. 14245–14254. Computer Vision Foundation / IEEE (2021)
7. Chen, C., Shi, X., Qin, Y., Li, X., Han, X., Yang, T., Guo, S.: Real-world blind super-resolution via feature matching with implicit high-resolution priors. In: Proceedings of the 30th ACM International Conference on Multimedia. pp. 1329–1338 (2022)
8. Chen, Z., Zhang, Y., Gu, J., Yuan, X., Kong, L., Chen, G., Yang, X.: Image super-resolution with text prompt diffusion. *CoRR* **abs/2311.14282** (2023)
9. Cheng, J., Liu, X., Zheng, K., Ke, P., Wang, H., Dong, Y., Tang, J., Huang, M.: Black-box prompt optimization: Aligning large language models without model training. arXiv preprint arXiv:2311.04155 (2023)
10. Dai, T., Cai, J., Zhang, Y., Xia, S.T., Zhang, L.: Second-order attention network for single image super-resolution. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11065–11074 (2019)
11. Ding, K., Ma, K., Wang, S., Simoncelli, E.P.: Image quality assessment: Unifying structure and texture similarity. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**(5), 2567–2581 (2022)
12. Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence* **38**(2), 295–307 (2015)
13. Dong, C., Loy, C.C., Tang, X.: Accelerating the super-resolution convolutional neural network. In: ECCV (2). Lecture Notes in Computer Science, vol. 9906, pp. 391–407. Springer (2016)
14. Fei, B., Lyu, Z., Pan, L., Zhang, J., Yang, W., Luo, T., Zhang, B., Dai, B.: Generative diffusion prior for unified image restoration and enhancement. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17–24, 2023. pp. 9935–9946 (2023)
15. Freeman, W.T., Pasztor, E.C., Carmichael, O.T.: Learning low-level vision. *International journal of computer vision* **40**, 25–47 (2000)
16. Gao, P., Han, J., Zhang, R., Lin, Z., Geng, S., Zhou, A., Zhang, W., Lu, P., He, C., Yue, X., et al.: Llama-adapter v2: Parameter-efficient visual instruction model. arXiv preprint arXiv:2304.15010 (2023)
17. Gu, J., Lu, H., Zuo, W., Dong, C.: Blind super-resolution with iterative kernel correction. In: CVPR. pp. 1604–1613. Computer Vision Foundation / IEEE (2019)

18. Gu, J., Shen, Y., Zhou, B.: Image processing using multi-code GAN prior. In: CVPR. pp. 3009–3018. Computer Vision Foundation / IEEE (2020)
19. Gu, S., Lugmayr, A., Danelljan, M., Fritsche, M., Lamour, J., Timofte, R.: DIV8K: diverse 8k resolution image dataset. In: ICCV Workshops. pp. 3512–3516. IEEE (2019)
20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778. IEEE Computer Society (2016)
21. Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross-attention control. In: ICLR. OpenReview.net (2023)
22. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: NIPS. pp. 6626–6637 (2017)
23. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
24. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NeurIPS (2020)
25. Ho, J., Salimans, T.: Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598* (2022)
26. Hosu, V., Lin, H., Szirányi, T., Saupe, D.: KonIQ-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Trans. Image Process.* **29**, 4041–4056 (2020)
27. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **43**(12), 4217–4228 (2021)
28. Kavar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. In: NeurIPS (2022)
29. Kavar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. *Advances in Neural Information Processing Systems* **35**, 23593–23606 (2022)
30. Kavar, B., Zada, S., Lang, O., Tov, O., Chang, H., Dekel, T., Mosseri, I., Irani, M.: Imagic: Text-based real image editing with diffusion models. In: CVPR. pp. 6007–6017. IEEE (2023)
31. Ke, J., Wang, Q., Wang, Y., Milanfar, P., Yang, F.: Musiq: Multi-scale image quality transformer. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5148–5157 (2021)
32. Kim, Y., Son, D.: Noise conditional flow model for learning the super-resolution. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops* (2021)
33. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597* (2023)
34. Li, J., Li, D., Xiong, C., Hoi, S.C.H.: BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. In: *ICML. Proceedings of Machine Learning Research*, vol. 162, pp. 12888–12900. PMLR (2022)
35. Li, W., Zhou, K., Qi, L., Lu, L., Lu, J.: Best-buddy gans for highly detailed image super-resolution. In: *AAAI*. pp. 1412–1420. AAAI Press (2022)
36. Liang, J., Zeng, H., Zhang, L.: Details or artifacts: A locally discriminative learning approach to realistic image super-resolution. In: CVPR. pp. 5647–5656. IEEE (2022)
37. Liang, J., Cao, J., Sun, G., Zhang, K., Gool, L.V., Timofte, R.: Swinir: Image restoration using swin transformer. In: *ICCVW*. pp. 1833–1844. IEEE (2021)

38. Lim, B., Son, S., Kim, H., Nah, S., Lee, K.M.: Enhanced deep residual networks for single image super-resolution. In: CVPR Workshops. pp. 1132–1140. IEEE Computer Society (2017)
39. 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. CoRR **abs/2308.15070** (2023)
40. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. *Advances in neural information processing systems* **36** (2024)
41. Liu, H., Wu, M., Yuan, K., Sun, M., Tang, Y., Zheng, C., Wen, X., Li, X.: Ada-dqa: Adaptive diverse quality-aware feature acquisition for video quality assessment. In: ACM Multimedia. pp. 6695–6704. ACM (2023)
42. Loshchilov, I., Hutter, F.: Fixing weight decay regularization in adam (2017)
43. Luo, F., Wu, X., Guo, Y.: And: Adversarial neural degradation for learning blind image super-resolution. *Advances in Neural Information Processing Systems* **36** (2024)
44. Menon, S., Damian, A., Hu, S., Ravi, N., Rudin, C.: PULSE: self-supervised photo upsampling via latent space exploration of generative models. In: CVPR. pp. 2434–2442. Computer Vision Foundation / IEEE (2020)
45. Michaeli, T., Irani, M.: Nonparametric blind super-resolution. In: ICCV. pp. 945–952. IEEE Computer Society (2013)
46. Mittal, A., Moorthy, A.K., Bovik, A.C.: No-reference image quality assessment in the spatial domain. *IEEE Trans. Image Process.* **21**(12), 4695–4708 (2012)
47. Mou, C., Wang, X., Xie, L., Zhang, J., Qi, Z., Shan, Y., Qie, X.: T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. CoRR **abs/2302.08453** (2023)
48. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: ICML. *Proceedings of Machine Learning Research*, vol. 139, pp. 8162–8171. PMLR (2021)
49. OpenAI: Gpt-4v(ision) system card. <https://openai.com/research/gpt-4v-system-card> (2023)
50. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
51. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* **1**(2), 3 (2022)
52. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
53. 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)
54. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems* **35**, 36479–36494 (2022)
55. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(4), 4713–4726 (2022)

56. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**(4), 4713–4726 (2023)
57. Sheikh, H.R., Sabir, M.F., Bovik, A.C.: A statistical evaluation of recent full reference image quality assessment algorithms. *IEEE Trans. Image Process.* **15**(11), 3440–3451 (2006)
58. Timofte, R., Agustsson, E., Van Gool, L., Yang, M.H., Zhang, L.: Ntire 2017 challenge on single image super-resolution: Methods and results. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. pp. 114–125 (2017)
59. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023)
60. Wang, J., Chan, K.C., Loy, C.C.: Exploring clip for assessing the look and feel of images. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 37, pp. 2555–2563 (2023)
61. Wang, J., Yue, Z., Zhou, S., Chan, K.C.K., Loy, C.C.: Exploiting diffusion prior for real-world image super-resolution. *CoRR* **abs/2305.07015** (2023)
62. Wang, J., Yue, Z., Zhou, S., Chan, K.C., Loy, C.C.: Exploiting diffusion prior for real-world image super-resolution. *arXiv preprint arXiv:2305.07015* (2023)
63. Wang, X., Xie, L., Dong, C., Shan, Y.: Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In: *ICCVW*. pp. 1905–1914. *IEEE* (2021)
64. Wang, X., Yu, K., Dong, C., Loy, C.C.: Recovering realistic texture in image super-resolution by deep spatial feature transform. In: *CVPR*. pp. 606–615. *Computer Vision Foundation / IEEE Computer Society* (2018)
65. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**(4), 600–612 (2004)
66. Wu, H., Zhang, Z., Zhang, E., Chen, C., Liao, L., Wang, A., Li, C., Sun, W., Yan, Q., Zhai, G., et al.: Q-bench: A benchmark for general-purpose foundation models on low-level vision. *arXiv preprint arXiv:2309.14181* (2023)
67. Wu, R., Yang, T., Sun, L., Zhang, Z., Li, S., Zhang, L.: Seesr: Towards semantics-aware real-world image super-resolution. *CoRR* **abs/2311.16518** (2023)
68. 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: *CVPR Workshops*. pp. 1190–1199. *IEEE* (2022)
69. Yang, T., Ren, P., Xie, X., Zhang, L.: Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. *arXiv preprint arXiv:2308.14469* (2023)
70. Yin, S., Fu, C., Zhao, S., Li, K., Sun, X., Xu, T., Chen, E.: A survey on multimodal large language models. *arXiv preprint arXiv:2306.13549* (2023)
71. You, Z., Li, Z., Gu, J., Yin, Z., Xue, T., Dong, C.: Depicting beyond scores: Advancing image quality assessment through multi-modal language models. *arXiv preprint arXiv:2312.08962* (2023)
72. Yuan, K., Kong, Z., Zheng, C., Sun, M., Wen, X.: Capturing co-existing distortions in user-generated content for no-reference video quality assessment. In: *ACM Multimedia*. pp. 1098–1107. *ACM* (2023)
73. Zhang, K., Liang, J., Gool, L.V., Timofte, R.: Designing a practical degradation model for deep blind image super-resolution. In: *ICCV*. pp. 4771–4780. *IEEE* (2021)

74. Zhang, K., Zuo, W., Zhang, L.: Learning a single convolutional super-resolution network for multiple degradations. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3262–3271 (2018)
75. Zhang, L., Zhang, L., Bovik, A.C.: A feature-enriched completely blind image quality evaluator. *IEEE Trans. Image Process.* **24**(8), 2579–2591 (2015)
76. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023)
77. Zhang, P., Wang, X.D.B., Cao, Y., Xu, C., Ouyang, L., Zhao, Z., Ding, S., Zhang, S., Duan, H., Yan, H., et al.: Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and composition. *arXiv preprint arXiv:2309.15112* (2023)
78. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR. pp. 586–595. Computer Vision Foundation / IEEE Computer Society (2018)
79. Zhang, R., Gu, J., Chen, H., Dong, C., Zhang, Y., Yang, W.: Crafting training degradation distribution for the accuracy-generalization trade-off in real-world super-resolution. In: ICML. Proceedings of Machine Learning Research, vol. 202, pp. 41078–41091. PMLR (2023)
80. Zhang, Y., Huang, X., Ma, J., Li, Z., Luo, Z., Xie, Y., Qin, Y., Luo, T., Li, Y., Liu, S., Guo, Y., Zhang, L.: Recognize anything: A strong image tagging model. *CoRR* **abs/2306.03514** (2023)
81. Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., Fu, Y.: Image super-resolution using very deep residual channel attention networks. In: ECCV (7). Lecture Notes in Computer Science, vol. 11211, pp. 294–310. Springer (2018)
82. Zhao, K., Yuan, K., Sun, M., Li, M., Wen, X.: Quality-aware pre-trained models for blind image quality assessment. In: CVPR. pp. 22302–22313 (2023)
83. Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592* (2023)