
Visual Autoregressive Modeling for Image Super-Resolution

Yunpeng Qu^{1,2} Kun Yuan² Jinhua Hao² Kai Zhao² Qizhi Xie^{1,2} Ming Sun² Chao Zhou²

Abstract

Image Super-Resolution (ISR) has seen significant progress with the introduction of remarkable generative models. However, challenges such as the trade-off issues between fidelity and realism, as well as computational complexity, have also posed limitations on their application. Building upon the tremendous success of autoregressive models in the language domain, we propose VARSR, a novel visual autoregressive modeling for ISR framework with the form of next-scale prediction. To effectively integrate and preserve semantic information in low-resolution images, we propose using prefix tokens to incorporate the condition. Scale-aligned Rotary Positional Encodings are introduced to capture spatial structures, and the Diffusion Refiner is utilized for modeling quantization residual loss to achieve pixel-level fidelity. Image-based Classifier-free Guidance is proposed to guide the generation of more realistic images. Furthermore, we collect large-scale data and design a training process to obtain robust generative priors. Quantitative and qualitative results show that VARSR is capable of generating high-fidelity and high-realism images with more efficiency than diffusion-based methods. Our codes are released at <https://github.com/quyp2000/VARSR>.

1. Introduction

Image Super-Resolution (ISR) aims to generate realistic high-resolution (HR) images from their degraded low-resolution (LR) counterparts. Traditional ISR methods focus on restoring LR images by assuming simple and known degradations (Dong et al., 2016; Zhang et al., 2018b), which limits their practicality in real-world scenarios with complex distortions (Gu et al., 2019a; Zhang et al., 2023b). Recent

methods leverage generative model priors to tackle ISR, with approaches based on GAN models (Zhang et al., 2021; Wang et al., 2021) and diffusion models (Ho et al., 2020). While these methods have achieved significant advancements, ISR, as an ill-posed problem, struggles to balance the realism and fidelity of the restoration results. GAN-based methods achieve higher fidelity metrics, but limitations in generation capability and the goal of over-fidelity optimization make it challenging to reproduce vivid and realistic textures (Wu et al., 2024a). Diffusion methods, based on their strong generative priors, can generate extremely rich image details. However, the random noise sampling approach and the gap between generative priors and LR distribution pose challenges for pixel-level fidelity (Yu et al., 2024). Building upon previous works, we aim to explore further existing frameworks to enhance fidelity and realism.

Introducing the generative capability of large models to enhance ISR is a prevalent trend. With the rich semantic priors of T2I models pretrained on massive datasets, many works (Yang et al., 2025; Wu et al., 2024b) have applied the powerful generative models to ISR tasks and achieved significant results. However, for ISR tasks demanding pixel-level fidelity, the diffusion model presents certain limitations, including the high computational complexity of iterative inference and potential issues with semantic hallucinations (Kim et al., 2025). With the success of autoregressive (AR) modeling in Large Language Models (LLMs) (Touvron et al., 2023; Achiam et al., 2023), visual autoregressive modeling (VAR) has gained great attention, exemplified by VQGAN (Esser et al., 2021) and DALL-E (Reddy et al., 2021), which apply discrete token prediction for image generation. VAR (Tian et al., 2024) takes a step further by quantizing images into scale-wise token maps and generating images through next-scale prediction, leading to excellent results across multiple generation tasks (Tang et al.; Zhang et al., 2024; Yao et al., 2024). This novel approach provides new insights into ISR, showcasing potential advantages over diffusion methods: (1) The coarse-to-fine next-scale prediction conforms to Markov unidirectional modeling, providing better structural preservation and naturally adapting to ISR tasks; (2) The reduced number of inference steps and lower complexity in preceding scales result in higher efficiency.

However, unlike other controllable generation tasks, ISR requires generation with pixel-level fidelity alongside seman-

¹Tsinghua University, Beijing, China ²Kuaishou Technology, Beijing, China. Correspondence to: Kun Yuan <yuankun03@kuaishou.com>.

tics reservation, presenting challenges in integrating VAR into ISR tasks: (1) How to efficiently and effectively incorporate LR condition information to generate high-fidelity images? (2) AR-based approaches map images into 1D token sequences for processing. To achieve enhanced spatial structure preservation in ISR, how to better represent the positional relationships between tokens? (3) AR-based approaches quantize continuous latent to discrete tokens for training with cross-entropy loss. How to avoid the loss of information caused by the quantization process to improve the pixel-level fidelity of ISR? (4) Optimization objectives for fidelity may impact the generation of rich image details. How to perceive low-quality distortions and guide the generation of more realistic images within the AR framework?

In this paper, we propose a novel *Visual Autoregressive modeling for image Super-Resolution (VARSR)* framework, with next-scale prediction. To enhance the integration of semantic information from LR images, we propose using *Prefix Tokens* as a proceeding scale to incorporate the LR condition as global guidance. To better capture the spatial structure between tokens, we utilize the *Scale-aligned Rotary Positional Encodings (SA-RoPE)* to calibrate the spatial positions of LR images and token maps. To predict the residual loss of the discrete quantizer, the *Diffusion Refiner* is introduced to model continuous distributions and achieve finer pixel-level restoration. To generate more realistic images, we define negative samples to learn low-quality distortions and propose *Image-based Classifier-free Guidance (CFG)* to guide the distribution for generating richer textures. Furthermore, to leverage powerful generative priors, we collect large-scale data and plan a comprehensive training process. Our **contributions** are as follows:

1. We first introduce VAR into the ISR field and propose VARSR, which is specifically designed to address LR conditions, spatial structure representation, quantization loss prediction, image-based CFG, and other issues to apply to ISR tasks.
2. We collect a large-scale, high-quality image dataset to establish a robust base model and design a training pipeline for fine-tuning downstream ISR tasks.
3. Through quantitative and qualitative analysis, VARSR demonstrates a strong capability in generating high-fidelity and high-realism images, achieving *the best performance* on multiple image quality metrics and $10\times$ *efficiency improvement* of diffusion methods.

2. Related work

Generative Priors for ISR. Existing ISR methods typically focus on blind recovery without assuming specific degradations (Huang et al., 2020; Bell-Kligler et al., 2019).

Generative priors are essential for addressing severe degradation in ISR tasks, capturing image structure and real-world distribution (Ho et al., 2020; Blattmann et al., 2022; Mou et al., 2024). While GAN-based methods show impressive results (Chen et al., 2022; Zhang et al., 2021; Wang et al., 2021), training stability and generative priors limit their ability to produce realistic details. Diffusion models are popular in ISR for strong generative capabilities (Saharia et al., 2022; Kawar et al., 2022; Qu et al., 2025), but may generate unrealistic details or hallucinations (Aithal et al., 2024; Narasimhaswamy et al., 2024). This may stem from their simplification of the Markov process, restricting the access to antecedent denoised trajectories (Gu et al., 2024).

Visual Autoregressive Models. Building upon the tremendous success of LLMs (Touvron et al., 2023; Achiam et al., 2023), visual autoregressive models utilize discrete quantizers such as VQVAE (Van Den Oord et al., 2017) to transform image patches into index-wise tokens, generating images based on next-token prediction (Yu et al., 2022; Wang et al., 2024b; Lee et al., 2022; Chang et al., 2022). However, the prediction of flattened tokens may lead to the loss of spatial structure. VAR (Tian et al., 2024) shifts from the next-token prediction to the next-scale prediction, significantly enhancing image generation quality and offering excellent scalability. VAR-based methods such as STAR (Ma et al., 2024), ControlAR (Li et al., 2024c) have expanded to other conditional generation tasks (*e.g.*, class-to-image (C2I), text-to-image (T2I)), yielding results comparable to diffusion models and validating the potential of VAR (Zhang et al., 2024; Yao et al., 2024; Roheda, 2024; Chen et al., 2024b).

3. Methods

3.1. Preliminary: Visual Autoregressive Modeling

Next-token Prediction is relied on by traditional AR models to generate patches at different positions in an image. Images are mapped into latent by a visual autoencoder and then tokenized into a series of token maps (x_1, x_2, \dots, x_T) using a discrete quantizer. AR model predicts the token x_T based on its preceding sequence $(x_1, x_2, \dots, x_{T-1})$ and condition c . The conditional probability can be expressed:

$$p(x_1, x_2, \dots, x_T) = \prod_{t=1}^T p(x_t | x_1, x_2, \dots, x_{t-1}, c), \quad (1)$$

Previous work (Tian et al., 2024) has pointed out that “next-token prediction” is insufficient for highly structured images, as it will disrupt the spatial structure and defy the unidirectional dependence assumption of autoregression.

Next-scale Prediction is reformulated based on these analyses (Tian et al., 2024), where the basic unit of autoregression becomes the token map of the entire scale. The visual en-

coder (e.g., VQVAE) first embeds the image I into a feature map $f \in \mathcal{R}^{h \times w \times d}$ and then quantizes the feature map f into K multi-scale token maps (r_1, r_2, \dots, r_K) .

$$\begin{aligned} f_k &= f - \sum_{m=1}^{k-1} \text{upsample}(\text{lookup}(V, r_m)), \\ r_k^{(i,j)} &= \arg \min_{v \in [V]} \left\| \text{lookup}(V, v) - f_k^{(i,j)} \right\|_2, \end{aligned} \quad (2)$$

where V is the VQVAE codebook, $r_k \in [V]^{h_k \times w_k}$ is the token map at scale k , and $\text{lookup}(V, v)$ retrieves vectors from codebook V based on index v . VAR predicts the next scale r_K based on previous outcomes and the condition:

$$p(r_1, r_2, \dots, r_K) = \prod_{k=1}^K p(r_k | r_1, r_2, \dots, r_{k-1}, c), \quad (3)$$

VAR with next-scale prediction advances AR models greatly. For ISR tasks demanding fidelity and realism, VAR’s progressive generation with coarse-to-fine coherence aligns with human perception and the Markov unidirectional assumption, ensuring high structural fidelity and aesthetic quality. Hence, applying VAR to ISR has broad prospects.

3.2. VARSR framework

We aim to enhance ISR via the VAR models. Unlike other controllable generation tasks (e.g., C2I, T2I), ISR necessitates maintaining semantic integrity and achieving precise pixel-level restoration. Therefore, when using the VAR, there still exist four critical issues requiring resolution:

- In what manner can the conditions derived from LR images be efficiently and effectively furnished to VAR to achieve high-fidelity semantic restoration?
- Tokens are concatenated into a 1D sequence in VAR, leading to the loss of spatial structures, which affects the fine-grained local restoration. How can the positional relationships of tokens be represented?
- VAR processes the image into discrete tokens through quantization, resulting in the significant loss of high-frequency details and impacting the pixel-level restoration. How to reduce the impact of quantization loss?
- The fidelity-oriented optimization objective may affect the generation of image details. How can we perceive distortions and quality factors in images to generate more realistic and higher-quality images?

To address these problems, we propose a framework called *Visual AutoRegressive modeling for image Super-Resolution (VARSR)*. As depicted in Fig. 1, it consists of

three main stages: conditional control generation, visual autoregression, and diffusion refinement. In the first stage, **to address the 1st problem**, we use an image encoder to map LR images to conditional tokens, which guide the generation of VAR by serving as *Prefix Tokens* (Sec. 3.2.1). In the second stage, a pretrained VAR model is employed. **To address the 2nd problem**, we implement *Scale-align rotary positional encoding (SA-RoPE)* in the transformer to calibrate the spatial positions of tokens at different scales (Sec. 3.2.2). In the third stage, **to address the 3rd problem**, a lightweight *Diffusion Refiner* estimates quantization residuals from generated discrete tokens to enhance image details (Sec. 3.2.3). **To address the 4th problem**, we further propose *Image-based classifier-free guidance (CFG)* as an additional control to learn distortions and generate better-quality images without additional training (Sec. 3.2.4).

3.2.1. CONDITIONAL CONTROL

The key challenge in ISR is how to effectively integrate LR information for controllable generation. Diffusion methods often use ControlNet (Zhang et al., 2023a) due to its superior performance over simpler approaches like concatenation or addition with noise. However, for AR models, ControlNet increases the computational burden and may lead to conflicts between LR priors and prefix-scale information.

Considering the next-scale prediction form of VAR, we adopt a more efficient way to introduce LR conditional control called *Prefix Tokens*. LR image I_{lr} is encoded through an image encoder \mathcal{E} to be mapped into conditional tokens $r_c = \mathcal{E}(I_{lr})$. r_c is used as the initial scale and placed at the start of all tokens in the form $(r_c, r_1, r_2, \dots, r_K)$, where r_c is of size $h \times w$ like the final scale r_K . VAR iteratively predicts the next scale by considering conditional tokens and the previous scales to effectively integrate LR priors.

In addition, HR images are categorized into two classes based on their quality during training. High- and low-quality categories are respectively associated with a positive embedding or a negative embedding for *quality-centered control* in VAR through a modulate layer. The objective is to impose *Image-based CFG*, as will be elaborated on in Sec. 3.2.4.

3.2.2. AUTOREGRESSIVE TRANSFORMER

As shown in Fig. 2, our VARSR utilizes a visual transformer for predicting multi-scale tokens. In training, the tokens r_k of the previous $K - 1$ scales are interpolated to next scale:

$$\hat{r}_k = \text{interpolate}(\text{lookup}(V, r_k)) \in \mathcal{R}^{h_{k+1} \times w_{k+1} \times d}, \quad (4)$$

Conditional tokens and interpolated tokens are concatenated as $(r_c, \hat{r}_1, \hat{r}_2, \dots, \hat{r}_{K-1})$ to predict the corresponding output $(\hat{r}_1, \hat{r}_2, \dots, \hat{r}_K)$ for each scale. A block-wise attention mask confines each r_k to focus solely on its preceding tokens $r \leq k$, with r_c accessible to all tokens to offer LR priors.

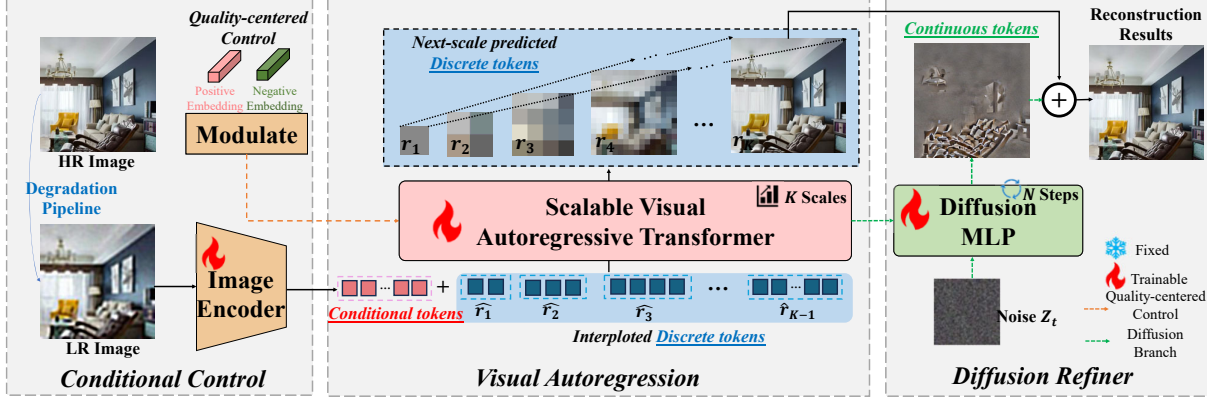


Figure 1. VARSR framework, which can be divided into three parts: (1) LR image is set as *Prefix Tokens* as condition. (2) VAR generates *discrete tokens* based on next-scale prediction. (3) Diffusion Refiner predicts the *continuous tokens* as quantization residuals.

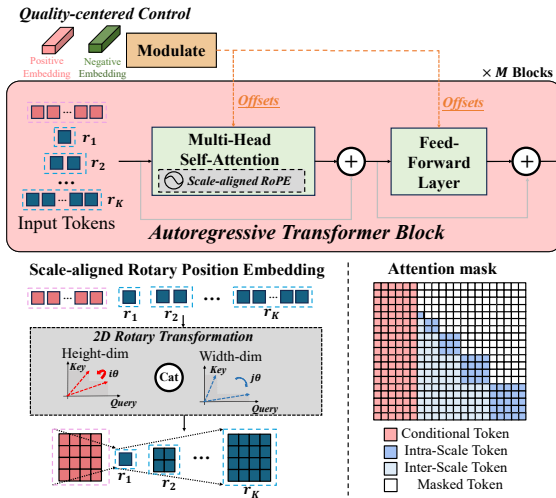


Figure 2. Internal structure of the autoregressive transformer. *SA-RoPE* represents the spatial structure. *Quality-centered control* generates offsets for autoregression.

Absolute positional encodings (APE) in transformers cannot capture 2D structure or represent the inter-scale relationships, which are crucial to incorporate spatial structure from LR images. Inspired by previous works using RoPE (Su et al., 2024) to represent the positions of token maps (Ma et al., 2024; Tang et al.), we propose an advancement with 2D RoPE to align spatial structures of LR images and multi-scale token maps, termed *Scale-aligned RoPE*. For the embedding $x_k^{(i,j)} \in \mathcal{R}^C$ of token $r_k^{(i,j)}$ in scale k at position (i, j) , we split the channels into two for representing 2D positions and normalize encodings for alignment:

$$\mathbf{R}_{\Theta}^C(x_k^{(i,j)}) = \begin{bmatrix} \mathbf{R}_{\Theta, (\frac{iH}{h_k})}^{\frac{C}{2}} & \mathbf{0}_{\frac{C}{2}} \\ \mathbf{0}_{\frac{C}{2}} & \mathbf{R}_{\Theta, (\frac{jW}{w_k})}^{\frac{C}{2}} \end{bmatrix}, \quad (5)$$

Where $\mathbf{R}_{\Theta}^C, \mathbf{0}_C \in \mathcal{R}^{C \times C}$ are the standard rotation matrix for RoPE and the zero matrices, respectively. Our *SA-RoPE* will be applied to the multiplication in self-attention to incorporate positional information for query and key of the two tokens x_m, x_n :

$$\mathbf{q}_m^T \mathbf{k}_n = (\mathbf{R}_{\Theta}^C(x_m) \mathbf{W}_q x_m)^T (\mathbf{R}_{\Theta}^C(x_n) \mathbf{W}_k x_n) \quad (6)$$

Where m, n represent their positional information, including the scale and position. It is worth noting that we also perform scale alignment on the LR condition r_c to incorporate the spatial structure of LR for next-scale restoration, enhancing the structural fidelity.

Lastly, the positive or negative embedding of the quality-centered control is introduced via the modulate layer, guiding the autoregression through the generation of offsets.

3.2.3. DIFFUSION REFINER

In Fig. 3, the image is quantized into discrete tokens, represented as indices of the codebook V . However, this discrete quantization introduces quantization loss, leading to the loss of high-frequency textures and restricting the upper bound of restoration, as VAR can only predict the quantized discrete vectors of the image. Such a loss is problematic for ISR, which demands precise pixel-level recovery. MAR (Li et al., 2024a) and HART (Tang et al.) utilize a diffusion loss to map predicted tokens to a continuous probability distribution, mitigating quantization loss. VARSR adopts the above approach of using diffusion as a refiner to convert predictions of categorical vector distribution into a continuous-valued space through a diffusion loss, thereby enhancing the upper bound of VAR’s capacity. The quantization residual is defined as continuous tokens z :

$$z = f - \sum_{k=1}^K \text{upsample}(\text{lookup}(V, r_k)), \quad (7)$$

We introduce a simple *Diffusion Refiner*, composed of a lightweight MLP (Li et al., 2024a). The *Diffusion Refiner* solely serves to map the probability distribution of quantized residuals, and does not have the capabilities of an additional ISR model. It leverages the final scale hidden states x_K as conditional control to generate a continuous distribution from noise z_t .

$$\mathcal{L}(x_K, z) = \mathbb{E}_{\varepsilon \sim \mathcal{N}(0,1), t} \left[\|\varepsilon - \varepsilon_\theta(z_t | t, x_K)\|^2 \right], \quad (8)$$

where ε is the sampled noise, t is the timestamp and $\varepsilon_\theta(\cdot)$ is the mapping function of the *Diffusion Refiner*. During inference, the sampled continuous tokens \hat{z} are combined with the predicted token maps for the final output.

3.2.4. IMAGE-BASED CLASSIFIER-FREE GUIDANCE

VAR, GANs, and Diffusion models all target fidelity as the optimization objective in ISR tasks, which may result in generated images being overly smooth and lacking in detail. It tends to retain distortions such as blur from LR images, leading to lower human-perceived quality.

Classifier-Free Guidance (CFG) (Ho & Salimans, 2022) has been widely used in T2I diffusion models, leveraging negative prompts to generate more realistic images by directed guiding of probabilistic distributions. Instead of simply setting the condition to null (Li et al., 2024d; Ho & Salimans, 2022) or using textual prompts, we aim to better perceive distortions and low-quality factors in images, hence proposing *Image-based CFG* to enhance generation quality. By learning low-quality image distributions during training, it allows us to guide the probability distribution during sampling towards generating higher-quality images, thereby expanding the upper bound.

In training, HR images are divided into high- and low-quality classes, with positive and negative embeddings c_p, c_n for control. In inference, to generate a higher-quality image I with condition r_c and c_p , the distribution is:

$$p(I|r_c, c_p) = \frac{p(c_p|r_c, I)p(I|r_c)}{p(c_p|r_c)}, \quad (9)$$

$$p(I|r_c) = p(c_p|r_c)p(I|r_c, c_p) + p(c_n|r_c)p(I|r_c, c_n),$$

Eq. 9 can be simplified to:

$$\begin{aligned} p(I|r_c, c_p) &= \frac{p(c_n|r_c)}{p(c_p|r_c)} \frac{p(c_p|r_c, I)p(I|r_c, c_n)}{1 - p(c_p|r_c, I)}, \\ \Rightarrow \nabla_I \log p(I|r_c, c_p) &= \nabla_I \log \frac{p(c_p|r_c, I)}{1 - p(c_p|r_c, I)} \\ &\quad + \nabla_I \log p(I|r_c, c_n), \end{aligned} \quad (10)$$

Furthermore, by applying the Bayesian formula, we can obtain:

$$\frac{p(c_p|r_c, I)}{1 - p(c_p|r_c, I)} = \frac{p(c_p|r_c)}{p(c_n|r_c)} \frac{p(I|r_c, c_p)}{p(I|r_c, c_n)}, \quad (11)$$

Substituting Eq. 11 into Eq. 10 yields Eq. 12. In CFG, a guiding scale λ_s balances diversity and realism.

$$\begin{aligned} \nabla_I \log p(I|r_c, c_p) &= \nabla_I \log p(I|r_c, c_n) \\ &\quad + \lambda_s (\nabla_I \log p(I|r_c, c_p) - \nabla_I \log p(I|r_c, c_n)), \end{aligned} \quad (12)$$

During inference, *Image-based CFG* are represented as:

$$\tilde{I} = \mathcal{F}(r_c, c_p) + \lambda_s (\mathcal{F}(r_c, c_p) - \mathcal{F}(r_c, c_n)). \quad (13)$$

where $\mathcal{F}(\cdot)$ is the mapping function of our VARSR.

3.3. Scaling up Database

Large-scale Data Collection. ISR methods require ample data to acquire generative priors and semantic understanding capabilities (Yu et al., 2024). Diffusion methods use Stable Diffusion (SD) (Rombach et al., 2022) trained on billions of text-image pairs for plentiful image priors. The open-source VAR base model falls short of our needs, as it could only generate 256×256 images and is limited in generated image quality, necessitating scaling high-quality training data for robust priors. Classical datasets often struggle to meet our high standards for image quality (e.g., ImageNet (Deng et al., 2009)) or quantity (e.g., DIV2K (Agustsson & Timofte, 2017), DIV8K (Gu et al., 2019b)). Therefore, we collect a ***new large-scale dataset with 4 million high-quality and high-resolution images across over 3k categories***, ensuring rich details and clear semantics.

Negative Samples. In Sec. 3.2.4, our VARSR utilizes negative embedding as inverse control to generate low-quality images. Therefore, we provide low-quality images as negative samples corresponding to negative embedding to learn low-quality distortions. We sample 50k low-quality images from various manually annotated image quality assessment (IQA) datasets (e.g., KonIQ10K (Hosu et al., 2020), CLIVE (Ghadiyaram & Bovik, 2016)) and image aesthetics assessment (IAA) dataset AVA (Murray et al., 2012) as negative samples added to our database.

Image Preprocess. Psychophysics research suggests that the richness of details in visual content impacts human perception of quality and aesthetics (He et al., 2014). Therefore, instead of the commonly-used random cropping, we resize images to 1.25 times the size for input and center-crop them during training. The aim is to enhance the comprehensive coverage of the foreground, which typically contains richer semantics and textures than the background, enabling VARSR to capture a broader range of visual attributes.

3.4. Training Procedure

Our entire training process is the same as the diffusion-based models and can be divided into a three-step process

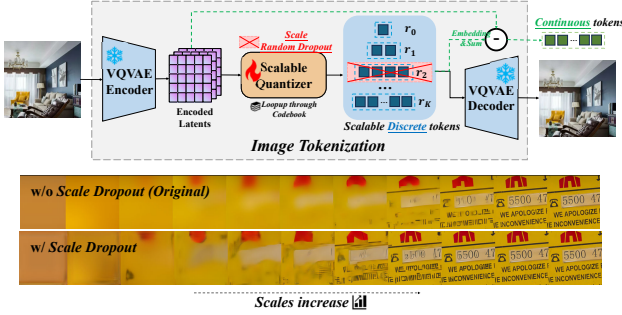


Figure 3. Image tokenization process of VQVAE. The quantizer converts the image latent to multi-scale *discrete tokens* while representing the quantization loss as *continuous tokens*.

for application in ISR tasks: (1) training the VAE, (2) pre-training on C2I/T2I tasks, and (3) finetuning on ISR. We train our model in all three stages to acquire a stronger generative prior for the ISR task.

VQVAE. In Fig. 3, the visualizations of the original VAR quantizer (Tian et al., 2024) reveal that image semantics are concentrated in the final few scales. This limits ISR, as earlier scale generations may lack information for later scales. Therefore, after the initial full training of the VQVAE, we added an additional training stage for the quantizer. Specifically, we follow the *scale random dropout* strategy in previous works (Li et al., 2024b; Kumar et al., 2024), where multi-scale quantized results are randomly discarded with a probability p_d . We freeze other parts and only train the quantizer. In Fig. 3, applying scale dropout preserves more semantic information in the earlier scales.

C2I Pretraining. We first train a powerful base VAR model on the C2I task to establish robust generation priors. Our training utilizes our large-scale dataset of over 3k categories, incorporating class information via start tokens and modulate layers in the transformer (Tian et al., 2024).

ISR Finetuning. We fine-tune the C2I pre-trained base model for the downstream ISR task to create the final VARSR. Both C2I Pretraining and ISR finetuning stages utilize the same loss with a coefficient λ to balance the cross-entropy loss of tokens and diffusion loss of the refiner.

4. Experiments

4.1. Experimental Setups

Datasets. We train VARSR on our large-scale dataset with negative samples using Real-ESRGAN’s degradation pipeline (Wang et al., 2021) to synthesize LR-HR image pairs. Both synthetic and real-world datasets are utilized




Figure 4. Limitations of current full-reference metrics (e.g., PSNR, SSIM, LPIPS). VARSR has generated images of higher perceptual quality for humans, yet it lags behind in certain metrics.

for a comprehensive evaluation. We create the synthetic validation set *DIV2K-VAL* by randomly cropping 3k patches from the DIV2K (Agustsson & Timofte, 2017) validation set, and for real-world evaluation, *DrealSR* (Cai et al., 2019) and *RealSR* (Wei et al., 2020) are center-cropped. Following (Wang et al., 2024a), all HR images have a resolution of 512×512 , and LR images are 128×128 .

Implemental Details. We train our VARSR following the procedure in Sec. 3.4, using a GPT-2 style (Radford et al., 2019) transformer with 24 blocks as the base model and a *Diffusion Refiner* with 6 blocks. We accelerate training by using pretrained VAR (Tian et al., 2024). We utilize an AdamW (Loshchilov & Hutter, 2017) optimizer with $\text{batchsize}=128$, $\text{weight decay}=5e-2$, and $\text{learning rate}=5e-5$. VQVAE, C2I pretraining, and ISR finetuning run for 10k, 40k, and 20k iterations, respectively. The loss balancing coefficient λ is 2.0, and the dropout ratio p_d is 0.1. The guidance scale λ_s linearly increases to 6.0 as the scale increases. Experiments are performed on 32 NVIDIA V100 GPUs.

Metrics. Reference-based metrics, including PSNR, SSIM (Wang et al., 2004) (Y channel), (Zhang et al., 2018a), and DISTS (Ding et al., 2022), are used for fidelity evaluation. FID (Heusel et al., 2017) measures the distribution distance between generated and reference images MANIQA (Yang et al., 2022), MUSIQ (Ke et al., 2021), and CLIPQA (Wang et al., 2023) are non-reference IQA metrics.

4.2. Comparison with SOTA

To verify the effectiveness, we compare our VARSR with other well-known SOTA GAN-based and Diffusion-based ISR methods¹, including BSRGAN (Zhang et al., 2021), Real-ESRGAN (Wang et al., 2021), SwinIR-GAN (Liang et al., 2021), DASR (Liang et al., 2022), LDM (Blattmann et al., 2022), StableSR (Wang et al., 2024a), DiffBIR (Lin et al., 2023), SinSR (Wang et al., 2024c), PASD (Yang et al., 2025) and SeeSR (Wu et al., 2024b).

¹All methods are tested based on their official code and models.

Table 1. Comparison with SOTA methods on synthetic and real-world benchmarks. Red and blue colors are the best and second-best.

Dataset	Metrics	GAN-based				Diffusion-based					AR-based	
		BSRGAN	Real-ESR	SwinIR	DASR	LDM	StableSR	DiffBIR	SinSR	PASD	SeeSR	VARSR
DIV2K-Val	PSNR↑	24.42	24.30	23.77	24.46	21.66	23.26	23.49	24.22	23.59	23.56	23.91
	SSIM↑	0.6164	0.6324	0.6186	0.6253	0.4752	0.5670	0.5568	0.5922	0.5899	0.5981	0.5980
	LPIPS↓	0.3511	0.3267	0.3910	0.3696	0.4887	0.3228	0.3638	0.3429	0.3611	0.3283	0.3260
	DISTS↓	0.2369	0.2245	0.2291	0.2533	0.2693	0.2116	0.2177	0.2157	0.2134	0.2008	0.2218
	FID↓	50.99	44.34	44.45	57.37	55.04	28.32	34.55	42.17	39.74	28.89	35.51
	MANIQA↑	0.3547	0.3756	0.3411	0.3104	0.3589	0.4173	0.4598	0.4101	0.4440	0.5046	0.5340
	CLIPQA↑	0.5253	0.5205	0.5213	0.4960	0.5570	0.6752	0.6731	0.6411	0.6573	0.6959	0.7347
	MUSIQ↑	60.18	59.76	57.21	53.96	57.46	65.19	65.57	61.46	66.58	68.35	71.27
	RealSR	PSNR↑	26.38	25.68	25.88	25.47	25.66	24.69	24.94	24.86	25.21	25.31
SSIM↑		0.7651	0.7614	0.7671	0.7575	0.6934	0.7090	0.6664	0.7191	0.7140	0.7284	0.7169
LPIPS↓		0.2656	0.2710	0.2614	0.3240	0.3367	0.3003	0.3485	0.3472	0.2986	0.2993	0.3504
DISTS↓		0.2124	0.2060	0.2061	0.2267	0.2324	0.2134	0.2257	0.2500	0.2125	0.2224	0.2470
FID↓		141.25	135.14	132.80	133.33	133.34	131.72	127.59	142.31	139.42	126.21	137.55
MANIQA↑		0.3763	0.3736	0.3561	0.2470	0.3375	0.4167	0.4378	0.3985	0.4418	0.5370	0.5570
CLIPQA↑		0.5114	0.4487	0.4433	0.3198	0.6053	0.6200	0.6396	0.6162	0.6009	0.6638	0.7006
MUSIQ↑		63.28	60.37	59.28	41.21	56.32	65.25	64.32	60.57	66.61	69.56	71.26
DRealSR		PSNR↑	28.70	28.61	28.20	29.75	27.78	27.87	26.57	28.26	27.45	28.13
	SSIM↑	0.8028	0.8052	0.7983	0.8262	0.7152	0.7427	0.6516	0.7443	0.7539	0.7711	0.7652
	LPIPS↓	0.2858	0.2819	0.2830	0.3099	0.3745	0.3333	0.4537	0.3743	0.3331	0.3142	0.3541
	DISTS↓	0.2144	0.2089	0.2103	0.2275	0.2417	0.2297	0.2724	0.2495	0.2322	0.2230	0.2526
	FID↓	155.62	147.66	146.38	155.36	164.87	148.18	160.67	173.42	173.40	147.00	155.87
	MANIQA↑	0.3441	0.3435	0.3311	0.2809	0.3342	0.3897	0.4602	0.3843	0.4551	0.5077	0.5362
	CLIPQA↑	0.5061	0.4525	0.4522	0.3813	0.5984	0.6321	0.6445	0.6302	0.6365	0.6893	0.7240
	MUSIQ↑	57.16	54.27	53.01	42.41	51.37	58.72	61.06	55.28	63.69	64.75	68.15

Quantitative Comparisons. The quantitative comparisons are shown in Tab. 1. *First, VARSR outperforms other SOTA methods by a wide margin in three no-reference IQA metrics:* MAINQA, CLIPQA, and MUSIQ, reflecting the capability to produce high-quality and realistic images. *Second,* in reference-based metrics (e.g., PSNR, SSIM), our VARSR approximates that of the best diffusion-based methods, yet still lags behind GAN methods. This is due to the trade-off between realism and fidelity, as VARSR and diffusion methods generate more textures and details, which may reduce fidelity metrics, especially for lower-quality HR images.

In Fig. 4, VARSR’s restoration exhibits higher quality in human perception, yet lags behind in certain reference-based metrics, as overly smoothed low-quality images tend to score better. This phenomenon highlights the limitations of current reference-based metrics, which have been confirmed in many studies (Wang et al., 2024a; Wu et al., 2024a). Theoretical derivations have verified the inherent contradiction between fidelity and quality (Blau & Michaeli, 2018).

Qualitative Comparisons. In Fig. 5, we visualize some of the ISR results from the test set. *First,* our VARSR consistently generates images with rich details and clear semantic information, showcasing the strong generative prior capability of autoregressive methods. *Second,* the image quality generated by GAN-based methods lags behind diffusion-based methods and our AR-based VARSR, highlighting the significance of generative priors. *Third, VARSR also outperforms diffusion-based methods in detail generation*

Table 2. Complexity comparison.

Methods	Params	Steps	Inference Time
StableSR	1409.1M	200	18.70s
PASD	1900.4M	20	6.07s
DiffBIR	1716.7M	50	5.85s
SeeSR	2283.7M	50	7.24s
VARSR(Ours)	1101.9M	10	0.59s

capability, aligning with the IQA metrics in Tab. 1. For example, in the 2nd and 4th row, VARSR distinguishes itself as **the sole method** that can correctly understand the semantics and generate houses and leaves with clear textures. In the 1st, 3rd, 5th, and 6th rows, VARSR demonstrates superior structural fidelity over other methods, generating images with well-defined structures and clear semantics. Additionally, VARSR excels in creating intricate animal fur, detailed leaves, and vivid textures. The results highlight VARSR’s robust ability to produce realistic and semantic-preserved images, even under severe degradation of LR images.

Complexity Analysis. Diffusion models have longer inference times due to multi-step noise sampling, while VARSR generates token maps at different scales with low iteration counts, ensuring minimal inference time. Additionally, with fewer token counts at early scales, the complexity of earlier steps is very low, unlike diffusion methods with consistent complexity per iteration. In Tab. 2, compared to diffusion models, **VARSR only requires 0.59s to generate an image, which is 10.1% of the second-ranked DiffBIR.** In addition, other optimization methods can also be applied to VARSR




Figure 5. Qualitative comparisons with different SOTA methods. **Zoom in for a better view.**




Figure 6. Effectiveness of the *LR Condition Mode*.

to further improve efficiency, such as using knowledge distillation to reduce the number of scales for inference.

4.3. Ablation Study

LR Condition. We use the form of *Prefix Tokens* to provide conditional control from LR images, comparing this mode with the other three LR condition modes: (1) Adding LR features directly in the transformer; (2) ControlNet (Zhang et al., 2023a) with a 1/2 depth; (2) ControlNet combined

with cross-attention (CA) (Yang et al., 2025). In Tab. 3, *Prefix Tokens* mode outperforms in all metrics. In Fig. 6, our method is the only one capable of generating the correct traffic light colors and clear text. These results showcase the effectiveness of the proposed *Prefix Tokens* pattern in extracting sufficient LR information for conditional guidance.

Scale-aligned RoPE. We validate the effectiveness of *SA-RoPE* by contrasting it with two scenarios: (1) Utilizing the original APE; (2) Applying *SA-RoPE* solely on discrete tokens. In Tab. 4, following the gradual application of *SA-RoPE*, fidelity metrics have shown a significant improvement. In Fig. 7, it is evident that applying *SA-RoPE* can better preserve the spatial structure, generating clear text and architectural details faithful to the original image.

Diffusion Refiner. We use the *Diffusion Refiner* to generate continuous tokens as the prediction of quantization residuals, supplementing the discrete tokens of the autoregressive transformer. In Tab. 5, after adding the *Diffusion Refiner*, both fidelity and image quality have significantly improved.

Table 3. Ablation on the *LR Condition Mode*.

Condition Mode	<i>DrealSR</i>					<i>RealSR</i>				
	SSIM↑	LPIPS↓	DISTS↓	MANIQA↑	MUSIQ↑	SSIM↑	LPIPS↓	DISTS↓	MANIQA↑	MUSIQ↑
Directly Add	0.6902	0.4011	0.3010	0.5188	66.46	0.6712	0.3742	0.2553	0.5221	69.90
ControlNet	0.7188	0.3824	0.2752	0.5260	66.78	0.7034	0.3683	0.2657	0.5298	69.70
ControlNet+CA	0.7314	0.3784	0.2630	0.5224	65.97	0.7084	0.3577	0.2552	0.5270	69.83
Prefix Tokens	0.7652	0.3541	0.2526	0.5362	68.15	0.7169	0.3504	0.2470	0.5570	71.26

 Table 4. Ablation on the *Scale-aligned RoPE*.

Exp	<i>DrealSR</i>			<i>RealSR</i>		
	✓	✗	✓	✗	✗	✓
Condition	✗	✗	✓	✗	✗	✓
Discrete	✗	✓	✓	✗	✗	✓
SSIM↑	0.7424	0.7603	0.7652	0.6992	0.7039	0.7169
LPIPS↓	0.3691	0.3645	0.3541	0.3654	0.3672	0.3504
DISTS↓	0.2617	0.2642	0.2526	0.2552	0.2576	0.2470
MANIQA↑	0.5238	0.5490	0.5362	0.5307	0.5555	0.5570
MUSIQ↑	68.14	68.41	68.15	71.08	71.32	71.26

 Table 5. Ablation on the *Diffusion Refiner*.

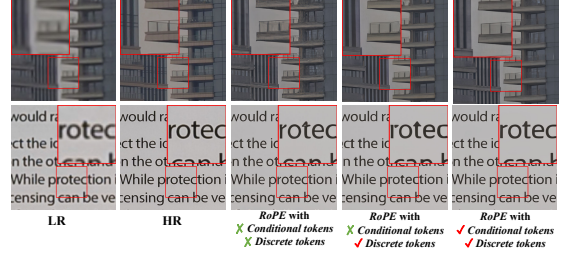


Metrics	<i>DrealSR</i>		<i>RealSR</i>	
	w/o Refiner	w/ Refiner	w/o Refiner	w/ Refiner
SSIM↑	0.7583	0.7652	0.7118	0.7169
LPIPS↓	0.3532	0.3541	0.3515	0.3504
DISTS↓	0.2539	0.2526	0.2503	0.2470
MANIQA↑	0.5297	0.5362	0.5399	0.5570
MUSIQ↑	67.99	68.15	70.74	71.26

 Table 6. Ablation on the *Image-based CFG*.

Metrics	<i>DrealSR</i>		<i>RealSR</i>	
	w/o CFG	w/ CFG	w/o CFG	w/ CFG
SSIM↑	0.8004	0.7652	0.7436	0.7169
LPIPS↓	0.2961	0.3541	0.2977	0.3504
DISTS↓	0.2161	0.2526	0.2103	0.2470
MANIQA↑	0.4052	0.5362	0.4326	0.5570
MUSIQ↑	59.70	67.63	66.32	71.26

In Fig. 8, we visualize discrete tokens and continuous tokens, showing that continuous tokens can effectively capture high-frequency details lost during quantization, resulting in images with richer semantic textures. A lightweight network suffices for the discrete-to-continuous mapping, and a larger refiner does not yield significant gains (Li et al., 2024a).

Image-based CFG. In Sec. 3.2.4, we introduce the *Image-based CFG* to generate more realistic and higher-quality images. Tab. 6 shows a notable enhancement in perceptual quality metrics, after incorporating the CFG. In Fig. 9, we show visualizations where images become clearer as the guidance scale λ_s increases. CFG results in significantly richer textures in the generated images, substantially improving perceptual quality to meet human preferences while maintaining correct semantics. However, excessive λ_s values may introduce artifacts not present in the original image. Therefore, our *Image-based CFG* achieves a balance between fidelity and realism by controlling λ_s . In practice,


 Figure 7. Effectiveness of the *Scale-aligned RoPE*.

 Figure 8. Visualization of *Continuous Tokens* and *Discrete Tokens*.

 Figure 9. Effect of the *Image-based CFG* with guidance scale λ_s .

$\lambda_s = 6.0$ is optimal, and it should increase linearly with scale for finer texture details in larger scales.

5. Conclusion

We explored VAR in ISR tasks and proposed the VARSR framework. To ensure pixel-level fidelity and realism, we made improvements in LR conditioning, structure representation, quantization prediction, and CFG guidance. We collected a large-scale dataset and designed a training process. Extensive experiments validate the performance of VARSR in generating high-fidelity and high-quality images.

Impact Statement

This paper introduces the visual autoregressive into the image super-resolution field. Therefore, any potential societal consequences or impacts related to ISR tasks apply here, as our work introduces new ideas that enhance ISR tasks with high efficiency and practicality.

References

- 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.
- Agustsson, E. and Timofte, R. NTIRE 2017 challenge on single image super-resolution: Dataset and study. In *CVPR Workshops*, pp. 1122–1131. IEEE Computer Society, 2017.
- Aithal, S. K., Maini, P., Lipton, Z., and Kolter, J. Z. Understanding hallucinations in diffusion models through mode interpolation. *Advances in Neural Information Processing Systems*, 37:134614–134644, 2024.
- Bell-Kligler, S., Shocher, A., and Irani, M. Blind super-resolution kernel estimation using an internal-gan. *Advances in Neural Information Processing Systems*, 32, 2019.
- Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10674–10685. IEEE, 2022.
- Blau, Y. and Michaeli, T. The perception-distortion trade-off. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6228–6237, 2018.
- Cai, J., Zeng, H., Yong, H., Cao, Z., and Zhang, L. Toward real-world single image super-resolution: A new benchmark and a new model. In *ICCV*, pp. 3086–3095. IEEE, 2019.
- Chang, H., Zhang, H., Jiang, L., Liu, C., and Freeman, W. T. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11315–11325, 2022.
- Chen, C., Shi, X., Qin, Y., Li, X., Han, X., Yang, T., and 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.
- Chen, C., Mo, J., Hou, J., Wu, H., Liao, L., Sun, W., Yan, Q., and Lin, W. Topiq: A top-down approach from semantics to distortions for image quality assessment. *IEEE Transactions on Image Processing*, 2024a.
- Chen, L., Wang, Z., Ren, S., Li, L., Zhao, H., Li, Y., Cai, Z., Guo, H., Zhang, L., Xiong, Y., et al. Next token prediction towards multimodal intelligence: A comprehensive survey. *arXiv preprint arXiv:2412.18619*, 2024b.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Ding, K., Ma, K., Wang, S., and Simoncelli, E. P. Image quality assessment: Unifying structure and texture similarity. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(5): 2567–2581, 2022.
- Dong, C., Loy, C. C., and Tang, X. Accelerating the super-resolution convolutional neural network. In *ECCV (2)*, volume 9906 of *Lecture Notes in Computer Science*, pp. 391–407. Springer, 2016.
- Esser, P., Rombach, R., and Ommer, B. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12873–12883, 2021.
- Gadre, S. Y., Ilharco, G., Fang, A., Hayase, J., Smyrnis, G., Nguyen, T., Marten, R., Wortsman, M., Ghosh, D., Zhang, J., Orgad, E., Entezari, R., Daras, G., Pratt, S. M., Ramanujan, V., Bitton, Y., Marathe, K., Musmann, S., Vencu, R., Cherti, M., Krishna, R., Koh, P. W., Saukh, O., Ratner, A. J., Song, S., Hajishirzi, H., Farhadi, A., Beaumont, R., Oh, S., Dimakis, A., Jitsev, J., Carmon, Y., Shankar, V., and Schmidt, L. Datacomp: In search of the next generation of multimodal datasets. In *NeurIPS*, 2023.
- Ghadiyaram, D. and Bovik, A. C. Massive online crowd-sourced study of subjective and objective picture quality. *IEEE Trans. Image Process.*, 25(1):372–387, 2016.
- Gu, J., Lu, H., Zuo, W., and Dong, C. Blind super-resolution with iterative kernel correction. In *CVPR*, pp. 1604–1613. Computer Vision Foundation / IEEE, 2019a.
- Gu, J., Wang, Y., Zhang, Y., Zhang, Q., Zhang, D., Jaitly, N., Susskind, J., and Zhai, S. Dart: Denoising autoregressive transformer for scalable text-to-image generation. *arXiv preprint arXiv:2410.08159*, 2024.
- Gu, S., Lugmayr, A., Danelljan, M., Fritsche, M., Lamour, J., and Timofte, R. DIV8K: diverse 8k resolution image dataset. In *ICCV Workshops*, pp. 3512–3516. IEEE, 2019b.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

- He, L., Gao, F., Hou, W., and Hao, L. Objective image quality assessment: a survey. *Int. J. Comput. Math.*, 91 (11):2374–2388, 2014.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, pp. 6626–6637, 2017.
- Ho, J. and Salimans, T. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.
- Hosu, V., Lin, H., Szirányi, T., and Saupe, D. Koniq-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Trans. Image Process.*, 29:4041–4056, 2020.
- Huang, Y., Li, S., Wang, L., Tan, T., et al. Unfolding the alternating optimization for blind super resolution. *Advances in Neural Information Processing Systems*, 33: 5632–5643, 2020.
- Kang, L., Ye, P., Li, Y., and Doermann, D. Convolutional neural networks for no-reference image quality assessment. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1733–1740, 2014.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4401–4410, 2019.
- Kawar, B., Elad, M., Ermon, S., and Song, J. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022.
- Ke, J., Wang, Q., Wang, Y., Milanfar, P., and Yang, F. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5148–5157, 2021.
- Kim, S., Jin, C., Diethe, T., Figini, M., Tregidgo, H. F., Mullokandov, A., Teare, P., and Alexander, D. C. Tackling structural hallucination in image translation with local diffusion. In *European Conference on Computer Vision*, pp. 87–103. Springer, 2025.
- Kumar, R., Seetharaman, P., Luebs, A., Kumar, I., and Kumar, K. High-fidelity audio compression with improved rvqgan. *Advances in Neural Information Processing Systems*, 36, 2024.
- Lee, D., Kim, C., Kim, S., Cho, M., and Han, W.-S. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11523–11532, 2022.
- Li, T., Tian, Y., Li, H., Deng, M., and He, K. Autoregressive image generation without vector quantization. *Advances in Neural Information Processing Systems*, 37:56424–56445, 2024a.
- Li, X., Qiu, K., Chen, H., Kuen, J., Gu, J., Raj, B., and Lin, Z. Imagefolder: Autoregressive image generation with folded tokens. *arXiv preprint arXiv:2410.01756*, 2024b.
- Li, X., Qiu, K., Chen, H., Kuen, J., Lin, Z., Singh, R., and Raj, B. Controlvar: Exploring controllable visual autoregressive modeling. *arXiv preprint arXiv:2406.09750*, 2024c.
- Li, Y., Zhang, K., Liang, J., Cao, J., Liu, C., Gong, R., Zhang, Y., Tang, H., Liu, Y., Demandolx, D., Ranjan, R., Timofte, R., and Gool, L. V. LSDIR: A large scale dataset for image restoration. In *CVPR Workshops*, pp. 1775–1787. IEEE, 2023.
- Li, Z., Cheng, T., Chen, S., Sun, P., Shen, H., Ran, L., Chen, X., Liu, W., and Wang, X. Controlar: Controllable image generation with autoregressive models. *arXiv preprint arXiv:2410.02705*, 2024d.
- Liang, J., Cao, J., Sun, G., Zhang, K., Gool, L. V., and Timofte, R. Swinir: Image restoration using swin transformer. In *ICCVW*, pp. 1833–1844. IEEE, 2021.
- Liang, J., Zeng, H., and Zhang, L. Efficient and degradation-adaptive network for real-world image super-resolution. In *European Conference on Computer Vision*, pp. 574–591. Springer, 2022.
- Lin, X., He, J., Chen, Z., Lyu, Z., Fei, B., Dai, B., Ouyang, W., Qiao, Y., and Dong, C. Diffbir: Towards blind image restoration with generative diffusion prior. *CoRR*, abs/2308.15070, 2023.
- Loshchilov, I. and Hutter, F. Fixing weight decay regularization in adam. 2017.
- Ma, X., Zhou, M., Liang, T., Bai, Y., Zhao, T., Chen, H., and Jin, Y. Star: Scale-wise text-to-image generation via auto-regressive representations. *arXiv preprint arXiv:2406.10797*, 2024.
- Mou, C., Wang, X., Xie, L., Wu, Y., Zhang, J., Qi, Z., and Shan, Y. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pp. 4296–4304, 2024.
- Murray, N., Marchesotti, L., and Perronnin, F. AVA: A large-scale database for aesthetic visual analysis. In *CVPR*, pp. 2408–2415. IEEE Computer Society, 2012.

- Narasimhaswamy, S., Bhattacharya, U., Chen, X., Dasgupta, I., Mitra, S., and Hoai, M. Handdiffuser: Text-to-image generation with realistic hand appearances. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2468–2479, 2024.
- Peebles, W. and Xie, S. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.
- Qu, Y., Yuan, K., Zhao, K., Xie, Q., Hao, J., Sun, M., and Zhou, C. Xpsr: Cross-modal priors for diffusion-based image super-resolution. In *European Conference on Computer Vision*, pp. 285–303. Springer, 2025.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. Learning transferable visual models from natural language supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021.
- Reddy, M. D. M., Basha, M. S. M., Hari, M. M. C., and Penchalaiah, M. N. Dall-e: Creating images from text. *UGC Care Group I Journal*, 8(14):71–75, 2021.
- Roheda, S. Cart: Compositional auto-regressive transformer for image generation. *arXiv preprint arXiv:2411.10180*, 2024.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and 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.
- Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D. J., and Norouzi, M. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4713–4726, 2022.
- Schuhmann, C. Clip+mlp aesthetic score predictor. <https://github.com/christophschuhmann/improved-aesthetic-predictor>, 2022.
- Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S., Crowson, K., Schmidt, L., Kaczmarczyk, R., and Jitsev, J. LAION-5B: an open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022.
- Shazeer, N. Fast transformer decoding: One write-head is all you need. *arXiv preprint arXiv:1911.02150*, 2019.
- Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- Su, S., Yan, Q., Zhu, Y., Zhang, C., Ge, X., Sun, J., and Zhang, Y. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3667–3676, 2020.
- Tang, H., Wu, Y., Yang, S., Xie, E., Chen, J., Chen, J., Zhang, Z., Cai, H., Lu, Y., and Han, S. Hart: Efficient visual generation with hybrid autoregressive transformer. In *The Thirteenth International Conference on Learning Representations*.
- Tian, K., Jiang, Y., Yuan, Z., Peng, B., and Wang, L. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *Advances in neural information processing systems*, 37:84839–84865, 2024.
- Timofte, R., Agustsson, E., Van Gool, L., Yang, M.-H., and 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.
- 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.
- Van Den Oord, A., Vinyals, O., et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- Wang, J., Chan, K. C., and Loy, C. C. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 2555–2563, 2023.
- Wang, J., Yue, Z., Zhou, S., Chan, K. C., and Loy, C. C. Exploiting diffusion prior for real-world image super-resolution. *International Journal of Computer Vision*, 132(12):5929–5949, 2024a.
- Wang, X., Xie, L., Dong, C., and Shan, Y. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *ICCVW*, pp. 1905–1914. IEEE, 2021.
- Wang, X., Zhang, X., Luo, Z., Sun, Q., Cui, Y., Wang, J., Zhang, F., Wang, Y., Li, Z., Yu, Q., et al. Emu3: Next-token prediction is all you need. *arXiv preprint arXiv:2409.18869*, 2024b.

- Wang, Y., Yang, W., Chen, X., Wang, Y., Guo, L., Chau, L.-P., Liu, Z., Qiao, Y., Kot, A. C., and Wen, B. Sinsr: diffusion-based image super-resolution in a single step. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 25796–25805, 2024c.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4): 600–612, 2004.
- Wei, P., Xie, Z., Lu, H., Zhan, Z., Ye, Q., Zuo, W., and Lin, L. Component divide-and-conquer for real-world image super-resolution. In *ECCV (8)*, volume 12353 of *Lecture Notes in Computer Science*, pp. 101–117. Springer, 2020.
- Wu, R., Yang, T., Sun, L., Zhang, Z., Li, S., and Zhang, L. Seesr: Towards semantics-aware real-world image super-resolution. In *CVPR*, pp. 25456–25467. IEEE, 2024a.
- Wu, R., Yang, T., Sun, L., Zhang, Z., Li, S., and Zhang, L. Seesr: Towards semantics-aware real-world image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 25456–25467, 2024b.
- Yang, Q., Chen, D., Tan, Z., Liu, Q., Chu, Q., Bao, J., Yuan, L., Hua, G., and Yu, N. HQ-50K: A large-scale, high-quality dataset for image restoration. *CoRR*, abs/2306.05390, 2023.
- Yang, S., Wu, T., Shi, S., Lao, S., Gong, Y., Cao, M., Wang, J., and Yang, Y. MANIQA: multi-dimension attention network for no-reference image quality assessment. In *CVPR Workshops*, pp. 1190–1199. IEEE, 2022.
- Yang, T., Wu, R., Ren, P., Xie, X., and Zhang, L. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. In *European Conference on Computer Vision*, pp. 74–91. Springer, 2025.
- Yao, Z., Li, J., Zhou, Y., Liu, Y., Jiang, X., Wang, C., Zheng, F., Zou, Y., and Li, L. Car: Controllable autoregressive modeling for visual generation. *arXiv preprint arXiv:2410.04671*, 2024.
- Yu, F., Gu, J., Li, Z., Hu, J., Kong, X., Wang, X., He, J., Qiao, Y., and Dong, C. Scaling up to excellence: Practicing model scaling for photo-realistic image restoration in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 25669–25680, 2024.
- Yu, J., Xu, Y., Koh, J. Y., Luong, T., Baid, G., Wang, Z., Vasudevan, V., Ku, A., Yang, Y., Ayan, B. K., Hutchinson, B., Han, W., Parekh, Z., Li, X., Zhang, H., Baldrige, J., and Wu, Y. Scaling autoregressive models for content-rich text-to-image generation. *Trans. Mach. Learn. Res.*, 2022, 2022.
- Zhai, X., Mustafa, B., Kolesnikov, A., and Beyer, L. Sigmoid loss for language image pre-training. In *ICCV*, pp. 11941–11952. IEEE, 2023.
- Zhang, K., Liang, J., Gool, L. V., and Timofte, R. Designing a practical degradation model for deep blind image super-resolution. In *ICCV*, pp. 4771–4780. IEEE, 2021.
- Zhang, L., Rao, A., and 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, 2023a.
- Zhang, Q., Dai, X., Yang, N., An, X., Feng, Z., and Ren, X. Var-clip: Text-to-image generator with visual autoregressive modeling. *arXiv preprint arXiv:2408.01181*, 2024.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pp. 586–595. Computer Vision Foundation / IEEE Computer Society, 2018a.
- Zhang, R., Gu, J., Chen, H., Dong, C., Zhang, Y., and Yang, W. Crafting training degradation distribution for the accuracy-generalization trade-off in real-world super-resolution. In *ICML*, volume 202 of *Proceedings of Machine Learning Research*, pp. 41078–41091. PMLR, 2023b.
- Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., and Fu, Y. Image super-resolution using very deep residual channel attention networks. In *ECCV (7)*, volume 11211 of *Lecture Notes in Computer Science*, pp. 294–310. Springer, 2018b. URL https://doi.org/10.1007/978-3-030-01234-2_18.

A. Implementation Details

A.1. Details of VARSR Framework

Table 7. Reconstruction results of the VQVAE on *RealSR* dataset.

Methods	Compression ratio	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	DISTS \downarrow	FID \downarrow	MANIQA \uparrow	CLIQQA \uparrow
VQVAE (w/ quantization)	16	34.33	0.9380	0.0428	0.0484	17.19	0.3432	0.4615
VQVAE (w/o quantization)	16	29.86	0.9006	0.0752	0.0688	34.27	0.3172	0.4327

VQVAE Tokenizer. We employ VQVAE (Van Den Oord et al., 2017) as the discrete tokenizer for VARSR, following the basic settings of VAR (Tian et al., 2024). The downsampling factor of VQVAE is $16\times$, and all scales share the same codebook with size $|V| = 4096$. Our VQVAE consists of two training stages: the initial full training and the application of the scale dropout strategy for the quantizer. Due to the *Diffusion Refiner* in our VARSR, which can predict quantization loss to some extent, we follow the alternating training strategy of HART (Tang et al.) for the initial full training. This involves incorporating continuous tokens for reconstruction with a probability of 50% (*i.e.*, bypassing the quantization process), ensuring that VQVAE can reconstruct images in both scenarios. The reconstruction results in both scenarios are illustrated in Tab. 7.

Image Encoder. We use a pyramid-style image encoder to extract information from degraded LR images and downsample it by a factor of 16 to obtain 32×32 feature space vectors. Therefore, the image encoder consists of 4 layers, each composed of two convolutional layers and an activation SiLU layer, including the last convolutional layer with a stride of 2×2 to reduce the size of the feature maps by half. All parameters of our image encoder are trainable.

Autoregressive Transformer. We utilize a Transformer based on the standard GPT-2 (Radford et al., 2019) architecture with a depth of 24 layers and a width of 1536. The modulate layer adopts the standard AdaLN (Peebles & Xie, 2023) form, which is added separately in each block to generate the offsets. The tokens are divided into 10 scales corresponding to resolutions ranging from 16×16 to 512×512 , with prefix tokens from LR added at the beginning. In the attention calculation, queries and keys reflect spatial positional information through our Scale-aligned RoPE.

During inference, KV-cache (Shazeer, 2019) is utilized to enhance the speed. Our VARSR supports the generation of higher-resolution images by concatenating them in a tiled manner, which is completely consistent with the diffusion-based method. Specifically, we uniformly divide the high-resolution image to be generated into overlapping grids. VARSR generates SR results for each grid separately and then tiles them together to obtain the complete image restoration result. We can batch process different grids in parallel to accelerate the inference process, significantly increasing computational speed.

Diffusion Refiner. Our *Diffusion Refiner* adheres to the specifications of (Li et al., 2024a), naturally supporting applications in CFG form. The *Diffusion Refiner* comprises a highly lightweight residual MLP as (He et al., 2016), with each residual block consisting of a LayerNorm, a linear layer, an activation SiLU layer, and another linear layer, merging with a residual connection. Our model consists of 6 layers, with a channel dimension of 1024. The hidden states from the autoregressive transformer serve as conditional controls and are introduced alongside timestamps T through the modulate module. Our noise schedule follows a cosine shape, with 1000 steps during training, and is resampled with 10 steps during inference, thus requiring minimal inference time. The parameter count of the lightweight *Diffusion Refiner* is less than 40M, accounting for 3% of the 1.1B model.

A.2. Details of Large-scale Dataset

Table 8. Comparison of our large-scale dataset with other well-known datasets.

Dataset	Scale	Resolution	MANIQA \uparrow	CLIQQA \uparrow	MUSIQ \uparrow	Aesthetic Score \uparrow
DIV2K	800	2k	0.43	0.62	71.15	5.57
Flickr2K	2650	2k	0.41	0.58	70.83	5.22
Large-scale data	Over 4M	2k	0.48	0.66	72.31	5.62

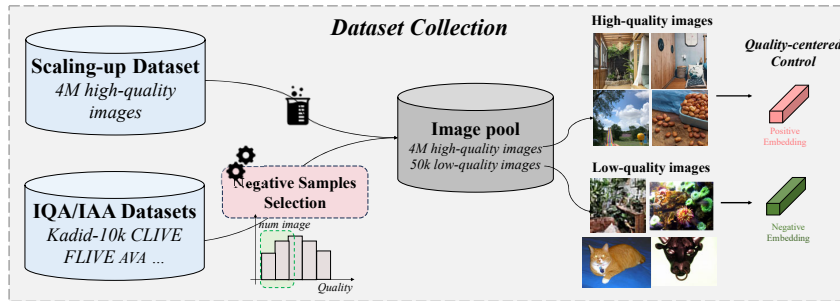


Figure 10. Database Collection. We collect a large-scale, high-quality dataset, and select 50k low-quality images from IQA and IAA datasets as negative samples, each with different *quality-centered controls* attached.

We collect billions of images from public datasets (e.g., LAION (Schuhmann et al., 2022), DataComp (Gadre et al., 2023)) and internal datasets. Employing a progressive filtering and semantic balance strategy, we have constructed a dataset comprising 4 million high-quality, high-resolution images.

Progressive Filtering. Images are progressively filtered using the following sub-metrics, with thresholds set for each: image metadata (including resolution, aspect ratio, and bits per pixel), IQA scores (comprising MANIQA (Yang et al., 2022), MUSIQ (Ke et al., 2021), CLIPIQA (Wang et al., 2023), and Aesthetic score (Schuhmann, 2022)), and texture richness (evaluated by the ratio of the power spectrum of high-frequency components (Yang et al., 2023) and the detection of blurred and flat regions (Li et al., 2023)). We establish the threshold for each metric at the top 70th percentile of its distribution within the DIV2K (Agustsson & Timofte, 2017) and Flickr2K (Timofte et al., 2017) datasets, ensuring that the images exhibit abundant details and high-frequency textures.

Table 9. Domain coverage of our large-scale dataset.

Categories	Indoor	Outdoor	Nature	Human	Plant	Object	Animal	Text	Food	Cartoon	Others
Rates	10%	9%	13%	28%	8%	9%	6%	4%	8%	2%	3%

Semantic Balance. To achieve diversity and balance of images across various domains, we implement semantic clustering based on the CLIP (Radford et al., 2021) and SigLIP (Zhai et al., 2023) models, and quantitatively select a sufficient number of images according to our predefined semantic categories, while specifically supplementing those categories with fewer images. As shown in the Tab. 9, we ensure that the dataset covers a wide range of category scenes and a relatively balanced proportion (Scenes with a broader semantic scope correspond to a higher proportion of images). The scenes include portraits, people, food, animals, natural landscapes, cartoons, cityscapes, indoor and outdoor scenes, thereby ensuring comprehensive coverage of visual concepts and richness in scene content.

B. Additional Experimental Results

B.1. Ability of Base Model

Table 10. Comparison of the base generative model on the C2I task.

Methods	Dataset	MANIQA ↑	CLIPIQA ↑	MUSIQ ↑
Original VAR	ImageNet	0.3250	0.5496	60.92
Original VAR	Large-scale data	0.4852	0.7015	72.74
VARSR	Large-scale data	0.5634	0.7314	74.32

To enhance the generative capabilities of the base model in the C2I task for downstream task migration, we pre-train the VAR model on our large-scale, high-quality dataset. In this section, we compare the capabilities of our VARSR base model



Figure 11. Visualization of C2I generation results from closely related categories.

trained on our large-scale dataset with the original VAR model (Tian et al., 2024) trained on ImageNet (Deng et al., 2009). As we are particularly interested in the model’s ability to generate detailed textures and ensure that it can still produce richly detailed images for downstream ISR task migration, we use non-reference IQA metrics for evaluation. In addition, it is unfair to use reference metrics such as FID for testing due to the differences in datasets.

Quantitative results are shown in Tab. 10, and in Fig. 11, we visualize some generated results from closely related categories. After using our high-quality large-scale dataset, the generated images show significant improvements in both objective metrics and subjective human observations. Furthermore, by incorporating our improvements (e.g., Scale-aligned RoPE, Diffusion Refiner) into the original VAR model, i.e., using the VARSR model, the quality further improved. This validates the effectiveness of our enhancements for the original C2I task. The results above validate the capabilities of our base model, which can generate images with rich details and clear semantics. It possesses a strong generative prior that can be leveraged by ISR tasks.

B.2. Comparisons with SOTA Methods

Table 11. Results of user study on real-world images.

Methods	BSRGAN	Real-ESRGAN	StableSR	PASD	SeeSR	VARSR(Ours)
Selection Rates	0.35%	0.7%	3.5%	17.3%	21.05%	57.1%

User Study. In order to comprehensively assess the performance of our VARSR in real-world scenarios, we conduct a user study on 100 randomly sampled LR real-world images from *DrealSR* (Cai et al., 2019) and *RealSR* (Wei et al., 2020). We compare our VARSR with five other GAN-based and diffusion-based ISR methods, including BSRGAN (Zhang et al., 2021), Real-ESRGAN (Wang et al., 2021), StableSR (Wang et al., 2024a), PASD (Yang et al., 2025), and SeeSR (Wu et al., 2024a). For each image, participants were presented with both the LR image and the restoration results of all ISR methods, and were then asked to indicate their choice for the best ISR result for the LR image. We invited 20 visual researchers to participate in the user study, and in total, we obtained 20×100 selection results. As shown in Tab. 11, our VARSR achieves **the highest selection rate of 57.1%**, far surpassing other methods, demonstrating the powerful capability of VARSR in real-world scenarios to generate realistic images that align with human aesthetics.

Comparisons on Real-world Images. To evaluate the performance of our VARSR in in-the-wild scenarios, we test different approaches on the *RealLR200* dataset (Wu et al., 2024a), which comprises 200 real-world images collected from previous studies (Lin et al., 2023; Wang et al., 2021) and from the internet. *RealLR200* dataset contains many highly degraded real-world images, such as historical photographs and extreme compression artifacts. Due to the absence of available reference HR images, we utilize only three no-reference IQA metrics: MANIQA (Yang et al., 2022), MUSIQ (Ke et al., 2021), and CLIPIQA (Wang et al., 2023).

Table 12. Quantitative comparison with SOTA methods on *RealLR200* dataset with no reference images. Red and blue colors are the best and second-best performance.

Dataset	Metrics	GAN-based			Diffusion-based					AR-based
		BSRGAN	Real-ESR	SwinIR	LDM	StableSR	DiffBIR	PASD	SeeSR	VARSR
<i>RealLR200</i>	MANIQA \uparrow	0.3671	0.3633	0.3741	0.3049	0.3688	0.4288	0.4295	0.4844	0.5177
	CLIPQA \uparrow	0.5698	0.5409	0.5596	0.5253	0.5935	0.6452	0.6325	0.6553	0.7513
	MUSIQ \uparrow	64.87	62.96	63.55	55.19	63.29	62.44	66.50	68.37	71.92

As shown in Tab. 12, our VARSR achieves **the best performance among all metrics**, which is consistent with the results on other synthetic and real-world datasets. In Fig. 15, we visualize some ISR results, demonstrating that VARSR can generate images with more realistic details compared to other methods. The results above validate the strong restoration capability of VARSR in real-world scenarios, showcasing its practical application value.

Table 13. Quantitative comparison with SOTA methods on other quality-based metrics. Red and blue colors are the best and second-best performance.

Dataset	Metrics	GAN-based			Diffusion-based					AR-based
		BSRGAN	Real-ESR	SwinIR	LDM	StableSR	DiffBIR	PASD	SeeSR	VARSR
<i>DIV2K-Val</i>	CNNIQA \uparrow	0.5492	0.5652	0.5402	0.5579	0.6274	0.6413	0.6269	0.6613	0.6661
	HyperIQA \uparrow	0.5682	0.5586	0.5235	0.5225	0.6100	0.6164	0.6158	0.6666	0.7031
	TOPIQ \uparrow	0.5413	0.5182	0.4796	0.4695	0.5923	0.6105	0.6165	0.6793	0.7020
<i>RealSR</i>	CNNIQA \uparrow	0.5513	0.5624	0.5281	0.5637	0.6029	0.6077	0.5938	0.6594	0.6692
	HyperIQA \uparrow	0.5617	0.5231	0.5093	0.4936	0.5703	0.5690	0.6001	0.6746	0.7038
	TOPIQ \uparrow	0.5502	0.5137	0.4882	0.4762	0.5579	0.5580	0.5920	0.6854	0.6991
<i>DRealSR</i>	CNNIQA \uparrow	0.4989	0.4849	0.5017	0.5367	0.5518	0.6025	0.5794	0.6132	0.6445
	HyperIQA \uparrow	0.5305	0.4938	0.5074	0.5050	0.5537	0.5992	0.6008	0.6583	0.6866
	TOPIQ \uparrow	0.5058	0.4622	0.4694	0.4807	0.5330	0.5831	0.5963	0.6534	0.6800

Additional IQA Metrics. In order to further validate VARSR’s ability to generate high-quality images that conform to real distributions, we select a broader range of IQA metrics to evaluate VARSR, including CNNIQA (Kang et al., 2014), HyperIQA (Su et al., 2020), and TOPIQ (Chen et al., 2024a). These IQA metrics were not used for the data filtering of our training dataset in 8. As shown in the Tab. 13, VARSR continues to achieve SOTA results, surpassing other methods. This validates the strong capability and robustness of VARSR to generate high-quality and content-rich images, rather than deceiving specific metrics.

Table 14. Quantitative comparison with SOTA methods with smaller finetuning data.

Metrics	<i>DrealSR</i>				<i>RealSR</i>			
	DiffBir	PASD	SeeSR	VARSR(LSDIR)	DiffBir	PASD	SeeSR	VARSR(LSDIR)
SSIM \uparrow	0.6516	0.7539	0.7711	0.7536	0.6664	0.7140	0.7284	0.7042
LPIPS \downarrow	0.4537	0.3331	0.3142	0.3716	0.3485	0.2986	0.2993	0.3614
DISTS \downarrow	0.2724	0.2322	0.2230	0.2620	0.2257	0.2125	0.2224	0.2540
MANIQA \uparrow	0.4602	0.4551	0.5077	0.5368	0.4602	0.4551	0.5077	0.5539
MUSIQ \uparrow	61.06	63.69	64.75	67.74	61.06	63.69	64.75	70.49

Finetuning on Smaller Datasets. The primary objective of using our large-scale data is for pretraining to acquire generative priors, similar to the powerful base models like Stable Diffusion (Rombach et al., 2022). Therefore, in this section, we present the results of finetuning VARSR on a smaller dataset after undergoing large-scale pretraining. We utilize the LSDIR (Li et al., 2023) dataset and 10k images from FFHQ (Karras et al., 2019) as high-quality data for finetuning ISR tasks, which is consistent with SeeSR, with a total of 95k high-quality images. In Tab. 14, when pretraining with our large-scale data and finetuning with smaller data, VARSR still performs exceptionally well, far surpassing other SOTA methods in perceptual

quality metrics. This is consistent with our conclusion when training on large-scale data, validating the superiority of the VARSR framework.


B.3. Ablation Study

Table 15. Ablation on the Training Database.


Metrics	<i>DrealSR</i>		<i>RealSR</i>	
	DIV2K	Large-scale Data	DIV2K	Large-scale Data
SSIM \uparrow	0.7750	0.7652	0.7284	0.7169
LPIPS \downarrow	0.3441	0.3541	0.3376	0.3504
DISTS \downarrow	0.2434	0.2526	0.2321	0.2470
MANIQA \uparrow	0.4885	0.5360	0.4990	0.5570
MUSIQ \uparrow	64.68	68.15	67.36	71.26

Table 16. Ablation on the Image Preprocess Strategy.

Metrics	<i>DrealSR</i>		<i>RealSR</i>	
	Random Crop	Resize&Crop	Random Crop	Resize&Crop
SSIM \uparrow	0.7772	0.7652	0.7304	0.7169
LPIPS \downarrow	0.3279	0.3541	0.3317	0.3504
DISTS \downarrow	0.2335	0.2526	0.2286	0.2470
MANIQA \uparrow	0.4577	0.5360	0.4639	0.5570
MUSIQ \uparrow	63.19	68.15	66.62	71.26



(a) Comparison of VARSR trained on *different dataset*.



(b) Comparison of VARSR trained on *image preprocess method*.

Figure 12. Comparison of *different dataset* and *image preprocess methods*.

Large-scale Dataset. In Tab. 15, we validate the importance of training VARSR with large-scale, high-quality data. We separately test the model trained on our large-scale dataset (over 4M images) and the model trained using the DIV2K (Agustsson & Timofte, 2017) series dataset (12k images, following the setting of (Yang et al., 2025)). The model trained on the DIV2K dataset utilized the original VAR (Tian et al., 2024) trained on ImageNet (Deng et al., 2009) as the pretrained model. Training with a large amount of high-quality data results in a significant performance improvement, as also confirmed by the visual results in Fig. 12(a). This demonstrates the effectiveness and necessity of using large-scale, high-quality data.

Image Preprocess. As depicted in Sec. 3.3, we preprocess the training images by first resizing and then cropping to include more foreground semantic information. In Tab. 16 and Fig. 12(b), we compare this preprocessing method with the

commonly used random cropping strategy. Our strategy achieves higher perceived quality, producing richer details, which demonstrates that high-frequency semantic information can be better preserved through our preprocessing.

Table 17. Ablation on the VQVAE Scale Dropout Strategy.

Metrics	DrealSR		RealSR	
	w/o dropout	$p_d = 0.1$	w/o dropout	$p_d = 0.1$
SSIM \uparrow	0.7567	0.7652	0.7155	0.7169
LPIPS \downarrow	0.3667	0.3541	0.3587	0.3504
DISTS \downarrow	0.2634	0.2526	0.2511	0.2470
MANIQA \uparrow	0.5410	0.5362	0.5414	0.5570
MUSIQ \uparrow	67.98	68.15	70.83	71.26




Figure 13. Effectiveness of the VQVAE Scale Dropout Strategy.

VQVAE Dropout Strategy. In Tab. 17, we compare the results of training the VARSR with the VQVAE tokenizer under two conditions, representing whether the VQVAE uses the *scale dropout strategy* during training with a dropout ratio $p_d = 0.1$. The tokenizer using the dropout strategy achieves significant improvements in all metrics, particularly in fidelity-based metrics, as the preceding scales can provide richer semantic information. As shown in Fig. 13, when the scale dropout strategy is not applied, all the fine texture details are generated in the final scale, which may result in some undesired artifacts in generated images. The condition with the dropout strategy applied can consistently generate images that are faithful to the original image.

B.4. Limitations



Figure 14. Limitations of VARSR. Due to limitations in the scope of training data coverage, it may be challenging to restore the correct semantics in certain extreme scenarios.

Although our VARSR base model is pretrained on the large-scale dataset of 4M images, the semantic coverage may still

have some limitations compared to the billions of data used by Stable Diffusion (Rombach et al., 2022). It is undeniable that some rare scenes may not be covered well. When the original image contains rare semantics and is severely damaged, it may not be possible to correctly restore the content. As shown in Fig. 14, the LR image’s severe distortion impedes the accurate restoration of the famous character Spider-Man in the left case, while the generated image of an ancient woman in the right case features simplistic and unrealistic head adornments.

C. More Visualization Results

In Fig. 16 and Fig. 17, we provide additional comparison visualization results with other methods, showcasing the robust capability of VARSR in generating high-fidelity and high-realistic images.

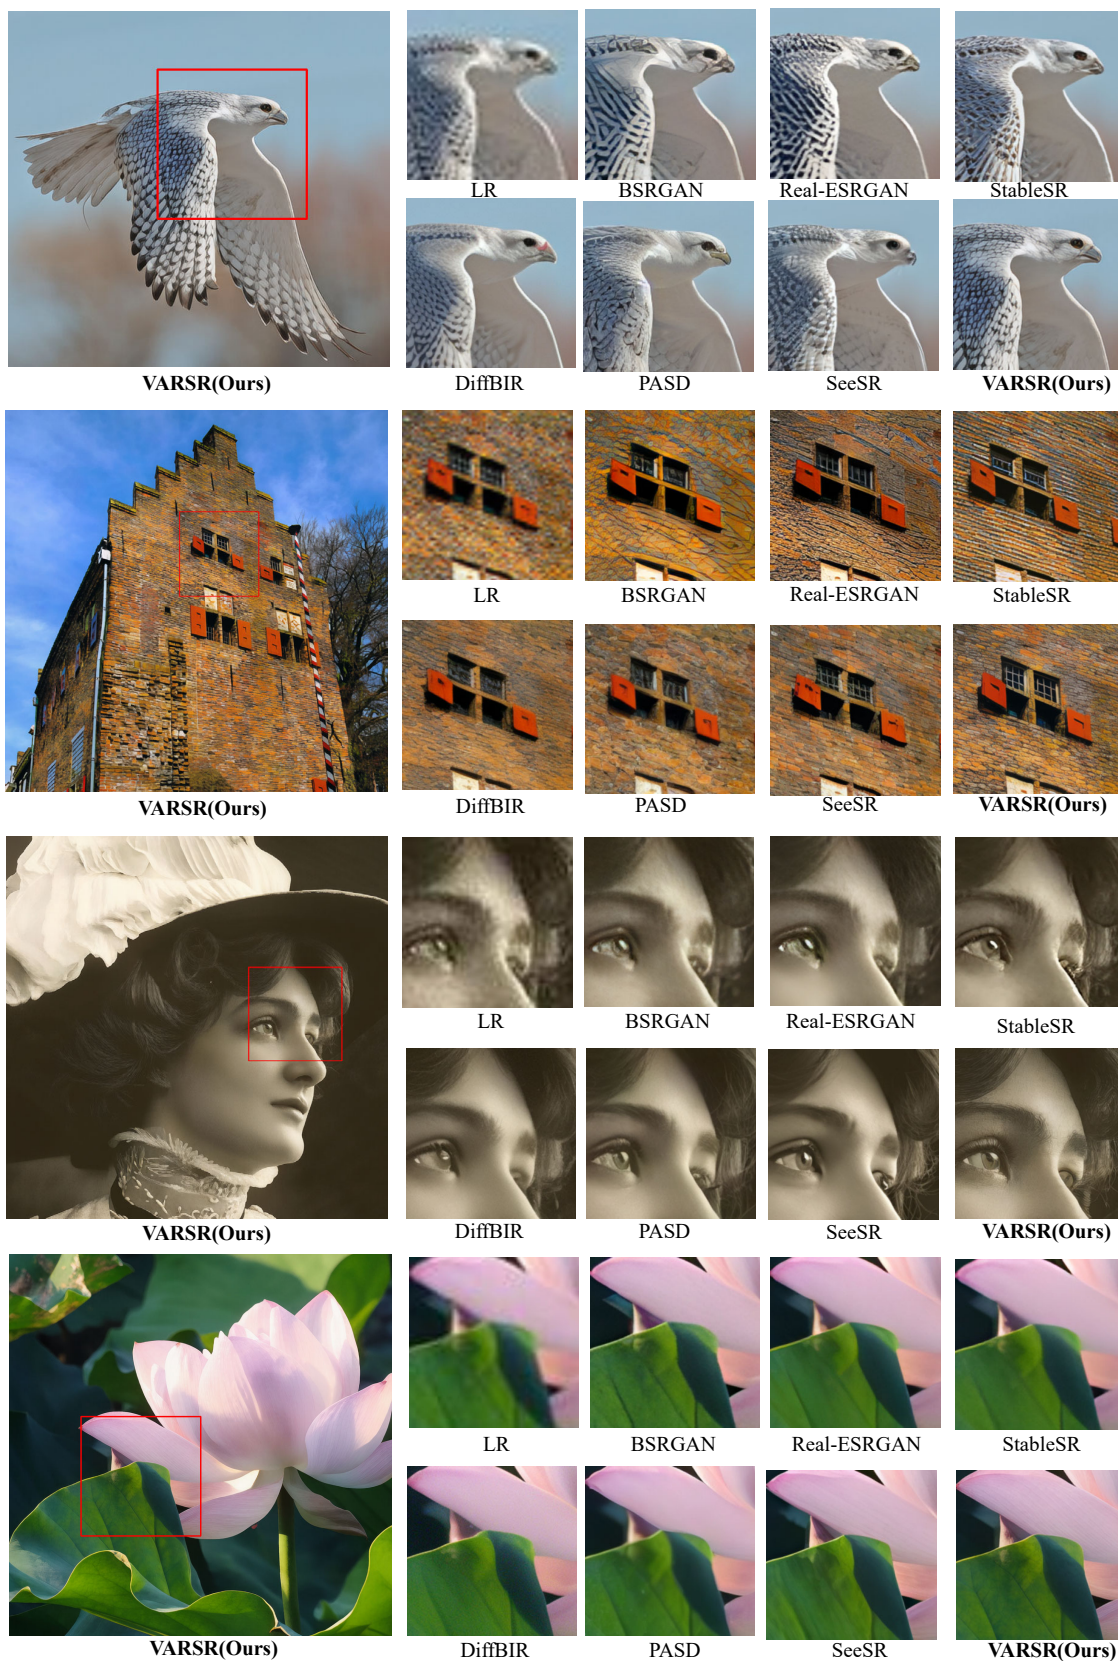


Figure 15. Comparisons with different methods on real-world images. **Zoom in for a better view.**

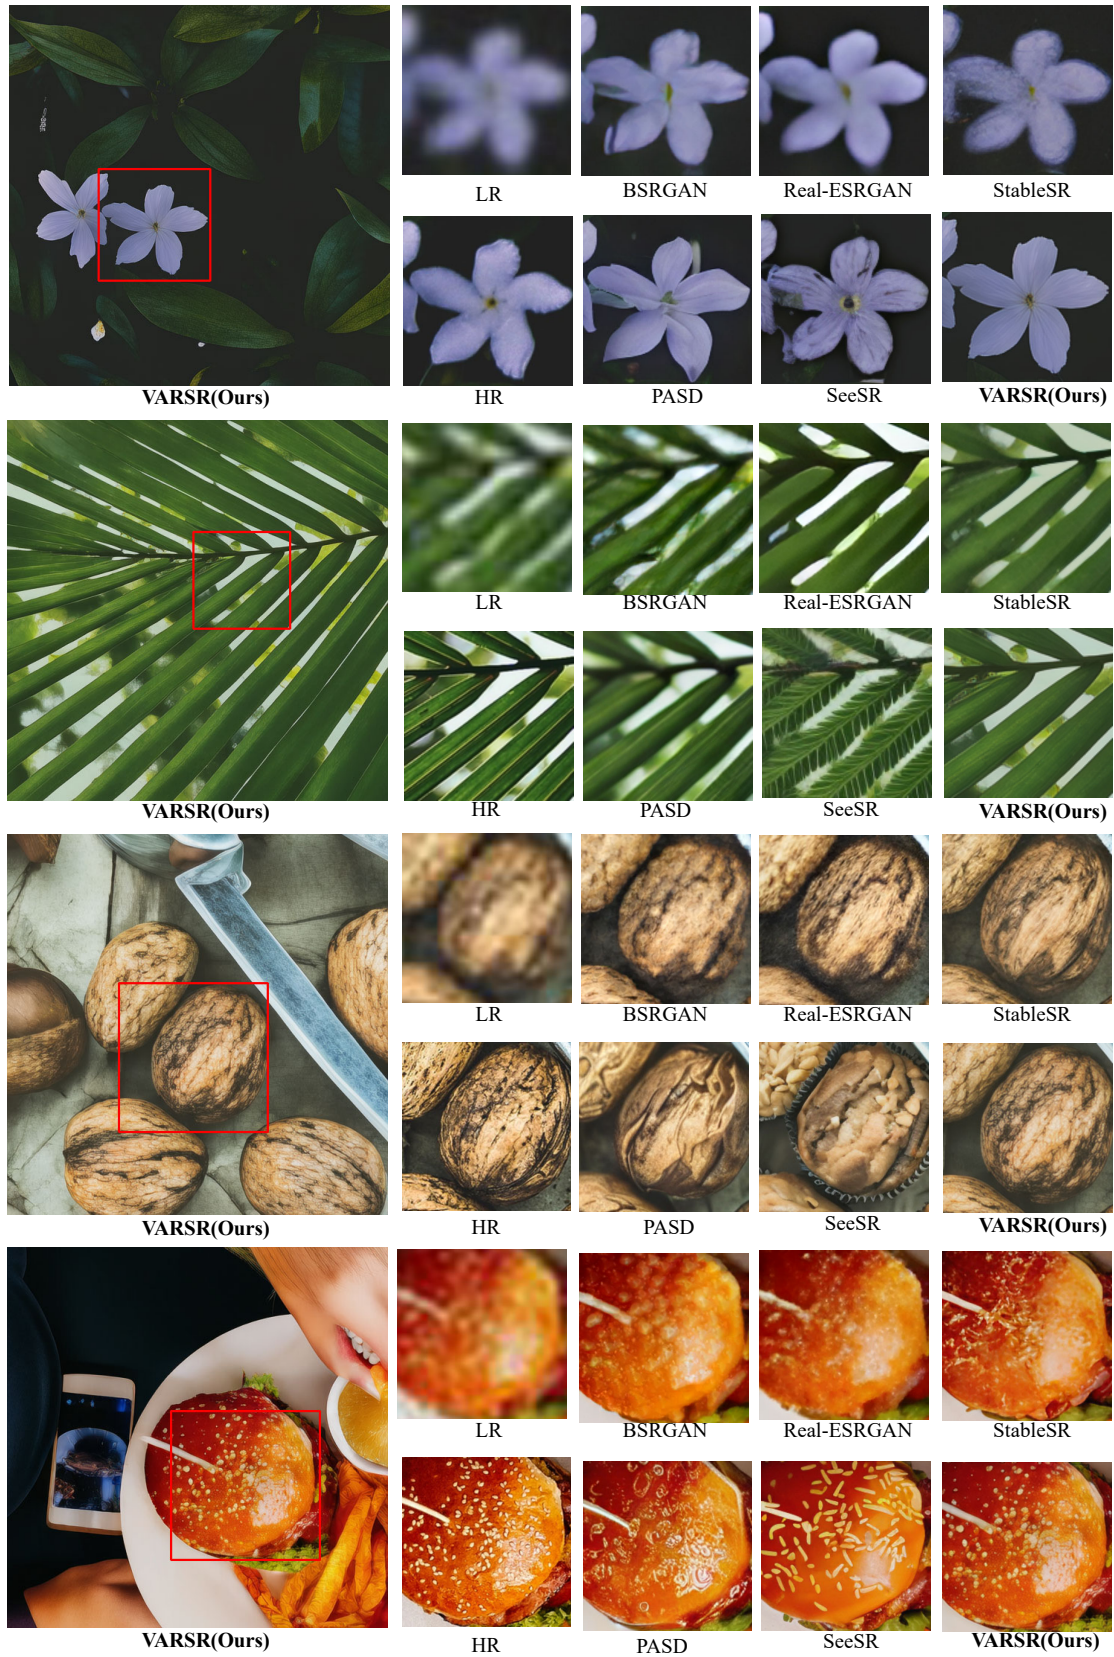


Figure 16. Additional qualitative comparisons with different SOTA methods (Part 1). **Zoom in for a better view.**

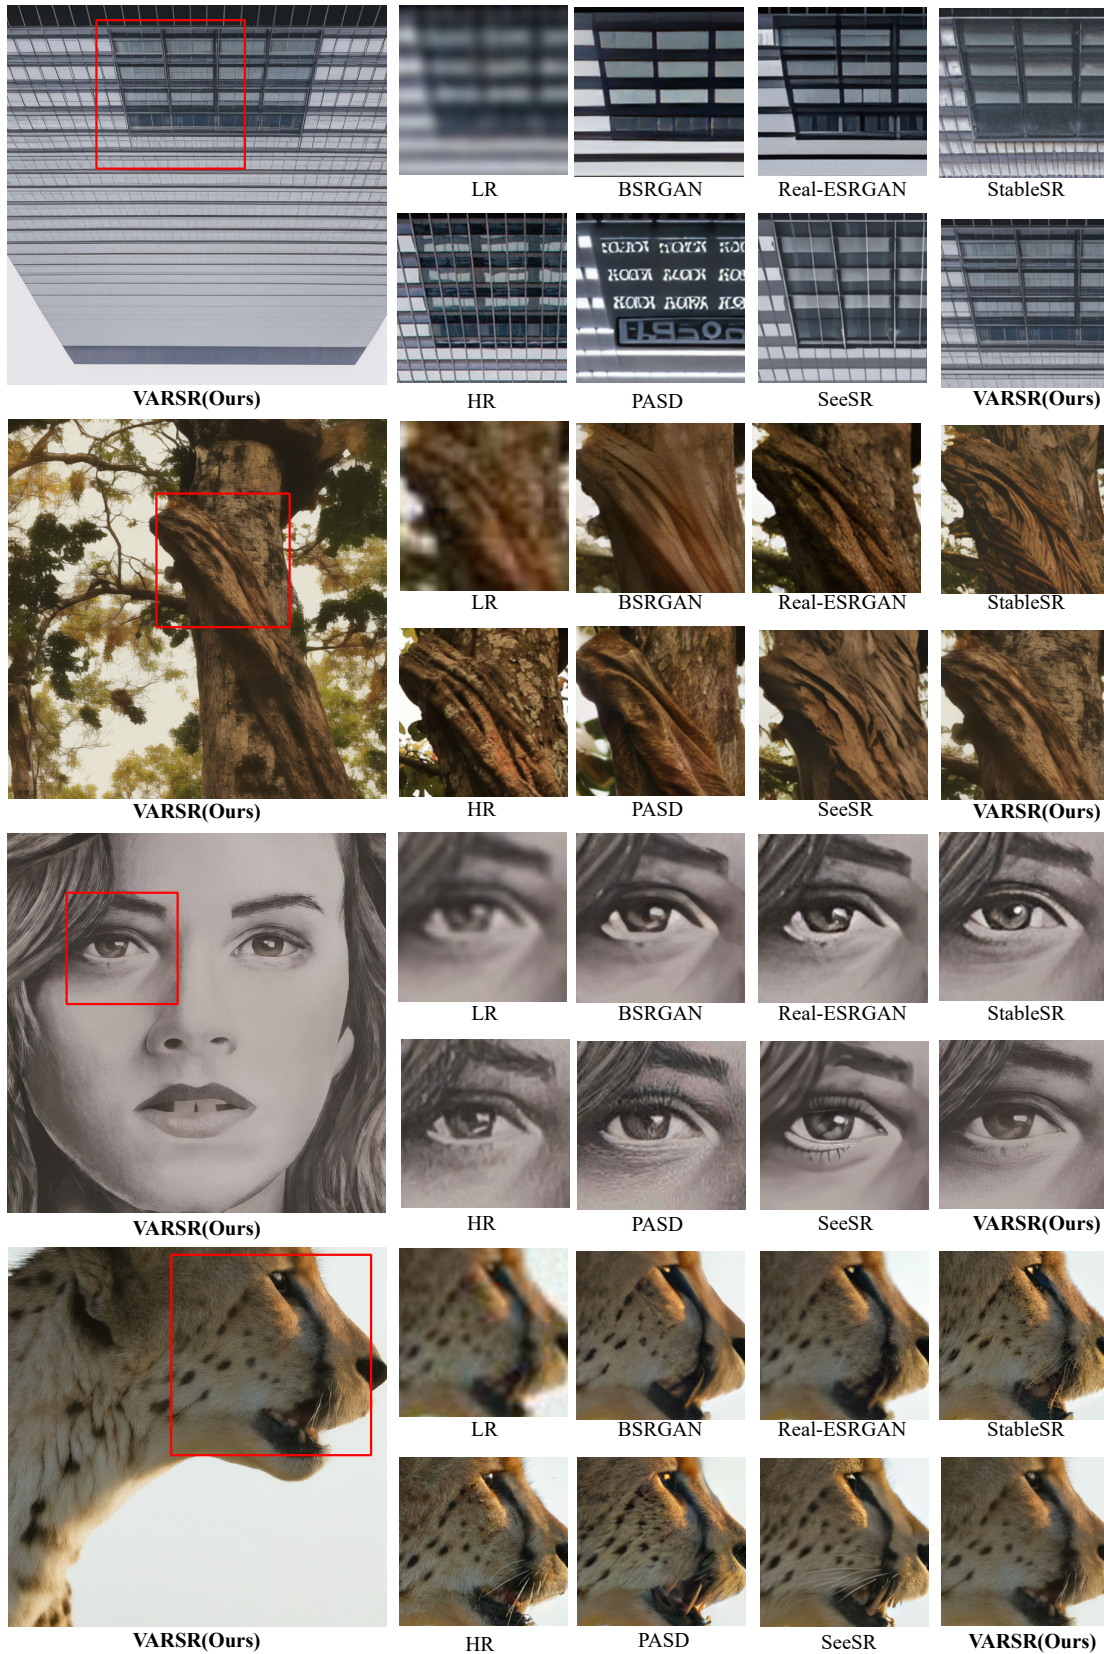


Figure 17. Additional qualitative comparisons with different SOTA methods (Part 2). Zoom in for a better view.