

Arbitrary-Scale 3D Gaussian Super-Resolution

Huimin Zeng^{1*}, Yue Bai¹, Yun Fu^{1,2}

¹Department of Electrical and Computer Engineering, Northeastern University

²Khoury College of Computer Science, Northeastern University
{zeng.huim, bai.yue, y.fu}@northeastern.edu

Abstract

Existing 3D Gaussian Splatting (3DGS) super-resolution methods typically perform high-resolution (HR) rendering of fixed scale factors, making them impractical for resource-limited scenarios. Directly rendering arbitrary-scale HR views with vanilla 3DGS introduces aliasing artifacts due to the lack of scale-aware rendering ability, while adding a post-processing upsampler for 3DGS complicates the framework and reduces rendering efficiency. To tackle these issues, we build an integrated framework that incorporates scale-aware rendering, generative prior-guided optimization, and progressive super-resolving to enable 3D Gaussian super-resolution of arbitrary scale factors with a single 3D model. Notably, our approach supports both integer and non-integer scale rendering to provide more flexibility. Extensive experiments demonstrate the effectiveness of our model in rendering high-quality arbitrary-scale HR views (**6.59 dB PSNR gain over 3DGS**) with a single model. It preserves structural consistency with LR views and across different scales, while maintaining real-time rendering speed (**85 FPS at 1080p**).

Code — <https://github.com/huimin-zeng/Arbi-3DGSR>

1 Introduction

High-resolution novel view synthesis (HRNVS) focuses on reconstructing 3D models from low-resolution (LR) sparse views and rendering high-resolution (HR) novel views. Such capability is critical for resource-limited scenarios such as low-bandwidth streaming (Grant et al. 2023; Solmaz and Van Gerven 2020) and AR/VR applications on portable devices (Fu et al. 2023), where only images with limited size can be transferred and efficiently processed. Previous Neural Radiance Field (NeRF) based methods (Yoon and Yoon 2023; Han et al. 2024; Huang et al. 2023; Wang et al. 2022; Huang et al. 2024) demonstrate powerful performance on HRNVS but suffer from rendering speed, making them impractical for resource-limited scenarios. Recent 3D Gaussian Splatting (3DGS) based methods (Shen et al. 2025; Feng et al. 2024; Xie et al. 2024; Ko et al. 2024; Bondarets et al. 2024; Xu et al. 2025; Yu et al. 2024a) utilize 3D Gaussian primitives to

describe scenes and achieve remarkable acceleration. However, existing 3DGS methods handle HRNVS at fixed integer scale factors (*e.g.*, $\times 2$ and $\times 4$) and require separate models for different scale factors, ignoring the intrinsic continuous characteristic of 3D world (Hu et al. 2024). They also fail to provide the flexibility for users to adjust rendering accuracy based on available resources. Therefore, a solution that enables 3D super-resolution of arbitrary scale factors using a single unified model is essential for practical HRNVS.

The most straightforward solution involves directly rendering views of target resolution using a scale-specific 3D model. However, this vanilla approach inevitably introduces aliasing artifacts (3DGS in Fig. 1), indicating its inflexibility and limited generalization ability for finer details. An alternative solution is to cascade novel view synthesis (NVS) models with an arbitrary-scale super-resolution (SR) upsampler to synthesize novel views of arbitrary scale factors. Yet, this solution complicates the framework and slows rendering due to the computational overhead of the upsamplers (see Sec. 5.4).

To overcome these limitations, we explore the task of arbitrary-scale 3D Gaussian super-resolution (Arbi-3DGSR), and propose an integrated framework that enables continuous rendering of HR views with arbitrary scale factors (including non-integer ones), using a single 3DGS model. We observe three challenges in Arbi-3DGSR: anti-aliasing NVS at various scale factors, constraining fine details of HR results without ground truth supervision, and maintaining structural consistency with LR views. Our framework tackles these challenges through three components: scale-aware rendering, generative prior-guided optimization and progressive super-resolving, respectively. Specifically, we approach Arbi-3DGSR with the following procedures: (1) **scale-aware rendering**: we utilize target scale factors to constrain the maximum signal frequency of the 3D model, and adapt the integration window size for accurate pixel shading; (2) **generative prior-guided optimization**: 2D generative priors from a diffusion model are leveraged to supervise fine details of rendered HR results. To avoid potential inconsistencies introduced by generative priors, we optimize with latent distillation and include orthogonal views for explicit supervision; and (3) **progressive super-resolving**: the training process is divided into multiple stages, gradually optimizing 3D Gaussian primitives to support larger scale factors while maintaining structural consistency.

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

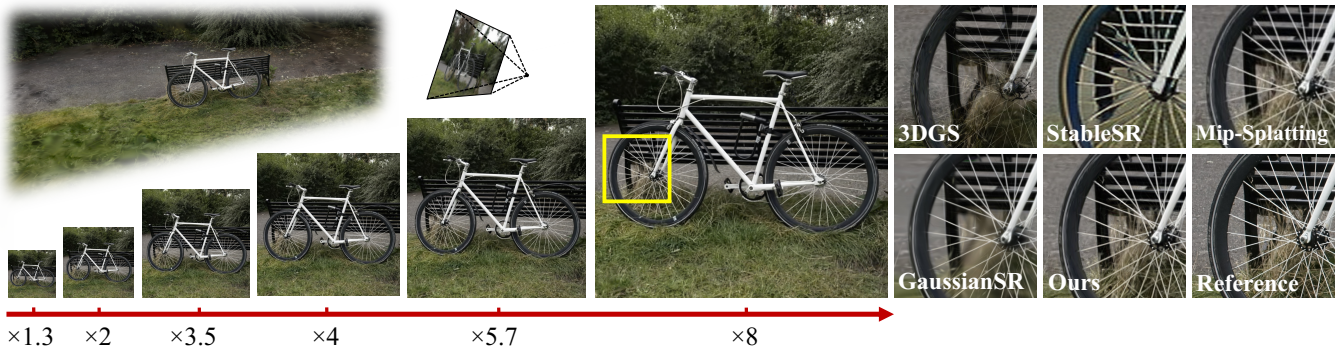


Figure 1: Visual results of typical solutions for arbitrary-scale 3D Gaussian super-resolution. Continuously rendering high-resolution novel views of different scale factors with vanilla 3DGS leads to aliasing artifacts. Cascaded solutions produce altered contents (*e.g.*, StableSR). Anti-aliasing Mip-Splatting and GaussianSR yield limited details.

Our contributions are summarized as follows:

- We make the first attempt to explore arbitrary-scale 3D Gaussian super-resolution (Arbi-3DGS), and propose an integrated framework that enables high-resolution rendering at varying scales using a single model.
- We introduce three key components: scale-aware rendering, generative prior-guided optimization, and progressive super-resolving to address critical challenges in anti-aliasing rendering, supervising fine details of HR results without ground truth, and maintaining structural consistency across various scale factors, respectively.
- Extensive experiments on four benchmarks show our superiority in rendering high-quality super-resolved results, including non-integer scale factors (*e.g.*, achieving a PSNR gain of **6.59 dB** over 3DGS at $\times 5.7$), while maintaining real-time speed of **85 FPS** for **1080p** rendering. Detailed ablations and visualizations provide more intuitions.

2 Related Work

2.1 High-Resolution Novel View Synthesis

Neural Radiance Field (NeRF) based methods (Yoon and Yoon 2023; Han et al. 2024; Lee, Li, and Lee 2024; Huang et al. 2023; Wang et al. 2022; Huang et al. 2024) enable NVS of arbitrary resolutions with implicit neural representations, but are computationally intensive due to complex architecture designs and volumetric rendering (Feng et al. 2024). While 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) provides efficient rendering of novel views. To provide high-resolution (HR) supervision, existing 3DGS-based methods (Feng et al. 2024; Xie et al. 2024; Bondarets et al. 2024; Shen et al. 2025; Feng et al. 2024; Xie et al. 2024; Ko et al. 2024) typically generate pseudo HR labels with pre-trained super-resolution models. DLGS (Xu et al. 2025) leverages a dual-lens system to provide geometric hints for HRNVS. Diffusion-guided methods (Yu et al. 2024a; Li et al. 2025a) introduce 2D generative priors and avoid explicit HR supervision. However, the aforementioned methods simply conduct HRNVS of fixed scale factors, ignoring the practical resource constraints and the continuous nature of the 3D world.

2.2 Anti-Aliasing 3D Gaussian-Splatting

Although 3DGS (Kerbl et al. 2023) enables real-time rendering, it suffers from aliasing artifacts under varying sampling rates (*i.e.*, different focal length and depth). Mip-Splatting (Yu et al. 2024b) introduces 3D smoothing filter and 2D Mip filter to constrain the maximal sampling frequency of Gaussian primitives. Analytic-Splatting (Liang et al. 2025) approximates Gaussian integrals over pixel areas and jointly models transmittance to enhance sampling robustness. SA-GS (Song et al. 2024) employs a frequency-aware 2D scale-adaptive filter to maintain consistent Gaussian distributions. MSGS (Yan et al. 2024) represent scenes with Gaussians of multiple scales, enabling adaptive selection for different sampling conditions to avoid aliasing artifacts. Unlike the HRNVS task that reconstructs from sparse LR views and renders novel HR views, anti-aliasing 3DGS inherently does not involve resolution changes, and the ground truth at different sampling rates is explicitly provided as supervision.

2.3 Arbitrary-Scale Super-Resolution

To enable continuous upscaling with a single network, MetaSR (Hu et al. 2019) dynamically adjusts filter weights based on scale factors. Zhao *et al.* (Zhao et al. 2024) enhance adaptability by fusing local features and scale factors. ArbiSR (Li et al. 2025b) employs a dual-level deformable implicit representation to address real-world degradations. COZ (Fu et al. 2024) boost degradation robustness by mixing features and coordinates of multiple points. BASR (Weng et al. 2024) model scaling with dual degradation representation to achieve cycle consistency. StableSR (Wang et al. 2024) introduces progressive aggregation sampling for resolution-agnostic generation. STAVSR (Shang et al. 2025) uses multi-scale priors to distinguish contents across different scales and locations. VideoINR (Chen et al. 2022) employs implicit neural representations to decode videos at arbitrary resolutions and frame rates. SAVSR (Li et al. 2024) introduces omni-dimensional scale-attention and bi-directional fusion to adapt across scales. Inspired by these works, we propose to incorporate scale factors into rendering process to improve the anti-aliasing ability across varying target resolutions.

3 Preliminaries

We provide a brief review of 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) and anti-aliasing filters introduced in Mip-Splatting (Yu et al. 2024b) to achieve aliasing-free novel view synthesis under varying focal lengths and depths.

3.1 3D Gaussian Splatting

3DGS models a scene with a set of explicit points $\{\mathbf{p}_i\}_{i=1}^N$ parameterized by opacity α , color \mathbf{c} , and 3D Gaussian primitive based geometry $\{G_i\}_{i=1}^N$. Each Gaussian primitive is described with a full 3D covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$ and center position $\mu \in \mathbb{R}^{3 \times 1}$ as follows:

$$G^{3D}(\mathbf{x}) = e^{-\frac{1}{2}(\mathbf{x}-\mu)^\top \Sigma^{-1}(\mathbf{x}-\mu)}, \quad (1)$$

where Σ is defined with a scaling matrix $S \in \mathbb{R}^{3 \times 1}$ and rotation matrix $R \in \mathbb{R}^{3 \times 3}$ as $RSS^\top R^\top$.

3.2 Anti-Aliasing Filtering

Aliasing occurs when a scene is reconstructed with fixed focal length and depth but rendered with lower sampling rates. Mip-Splatting (Yu et al. 2024b) band-limits each Gaussian primitive through a two-stage filtering, such that its highest signal frequency remains below the Nyquist frequency determined by input views.

3D Smoothing Filter. Projecting a continuous 3D signal onto the 2D screen plane is a sampling operation. According to Nyquist–Shannon sampling theorem (Nyquist 1928; Shannon 1949), given discrete samples taken at frequency \hat{r} , reconstructed continuous signals can theoretically yield frequency r up to $\frac{\hat{r}}{2}$ (*i.e.*, $\hat{r} \geq 2r$). Thus, considering a Gaussian primitive G_i^{3D} observed in K views with focal lengths $\{f_k\}_{k=1}^K$ in pixel units and depths $\{d_k\}_{k=1}^K$ in 3D world space, its maximum sampling rate is:

$$\hat{r}_i = \max \left(\left\{ \mathbb{I}_k(G_i^{3D}) \cdot \frac{f_k}{d_k} \right\}_{k=1}^K \right), \quad (2)$$

where $\mathbb{I}_k(G_i^{3D})$ evaluates the visibility of G_i^{3D} . To cap the highest frequency of reconstructed 3D model, Mip-Splatting imposes a low-pass Gaussian filter on each 3D Gaussian primitive G_i^{3D} , with hypermeter γ controlling the filter size:

$$G_i^{3D}(\mathbf{x})_{mip} = \sqrt{\frac{|\Sigma_i|}{|\Sigma_i + \frac{\gamma}{\hat{r}_i} \mathbf{I}}}} e^{-\frac{1}{2}(\mathbf{x}-\mu_i)^\top (\Sigma_i + \frac{\gamma}{\hat{r}_i} \mathbf{I})^{-1}(\mathbf{x}-\mu_i)}. \quad (3)$$

2D Mip Filter. During rendering, each 3D Gaussian in world space is projected to image plane as a 2D Gaussian, characterized by position $\hat{\mu} = \mathbf{P}\mathbf{W}[\mu, 1]^\top$ and covariance $\hat{\Sigma} = \mathbf{J}\mathbf{W}\Sigma\mathbf{W}^\top\mathbf{J}^\top$. Here \mathbf{P} , \mathbf{W} and \mathbf{J} denote the projection matrix, extrinsic matrix, and Jacobian of projective transformation, respectively. When shading a pixel $\hat{\mathbf{x}}$, Mip-Splatting imposes a 2D Gaussian filter over 2D Gaussians, thereby providing an integration window area and effectively preventing 2D Gaussian signals from being too small:

$$G_i^{2D}(\hat{\mathbf{x}})_{mip} = \sqrt{\frac{|\hat{\Sigma}_i|}{|\hat{\Sigma}_i + \varepsilon \mathbf{I}}}} e^{-\frac{1}{2}(\hat{\mathbf{x}}-\hat{\mu}_i)^\top (\hat{\Sigma}_i + \varepsilon \mathbf{I})^{-1}(\hat{\mathbf{x}}-\hat{\mu}_i)}, \quad (4)$$

where ε is a hypermeter to cover a single pixel. The color of pixel $\hat{\mathbf{x}}$ is then accumulated as follows:

$$\mathbf{C}(\hat{\mathbf{x}}) = \sum_{i=1}^N \alpha_i c_i G_i^{2D}(\hat{\mathbf{x}})_{mip} \prod_{j=1}^{i-1} (1 - \alpha_j G_j^{2D}(\hat{\mathbf{x}})_{mip}), \quad (5)$$

where α_i and c_i are the opacity and color of the i -th primitive.

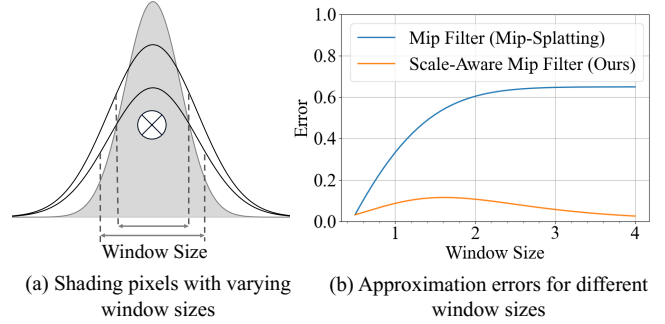


Figure 2: (a) Accurate pixel shading requires aligning the integration window with the actual pixel size. (b) Approximation error analysis regarding the window size, where the proposed method results in low approximation errors.

4 Method

Given a set of low-resolution (LR) views, our goal is to reconstruct a 3D scene and render high-resolution (HR) views at a target scale factor s (*i.e.*, the enlarging ratio from LR to HR views). We introduce a unified framework that comprises three key components: scale-aware rendering, generative prior-guided optimization, and progressive super-resolving. Scale-aware rendering is applied in both training and inference, enabling aliasing-free HRNVS at target scale factor s . During training, generative prior-guided optimization constrains fine details in the absence of ground truth, while progressive super-resolving gradually increases target scales to preserve structural consistency across varying scale factors.

4.1 Scale-Aware Rendering

The highest frequency of the reconstructed 3D Gaussian signal is fixed, whereas changing output resolution alters image-plane sampling density. A lower sampling rate can fall below the Nyquist frequency and cause aliasing, while a higher rate potentially blurs detail. Adjusting Gaussian bandwidth and per-pixel integration window both help to align the Gaussian signal with the pixel area. Therefore, we propose the following scale-aware filterings.

3D Scale-Aware Smoothing Filter. The maximum sampling rate of a Gaussian primitive is jointly determined by pixel density ρ , focal length f , and camera depth d with $\hat{r} = \frac{f\rho}{d}$. Hence, $\hat{r}(s) = \frac{f\rho \cdot s}{d}$ when the output resolution is scalable with factor s . For a Gaussian primitive G_i^{3D} that is visible in K views, its maximum sampling rate is:

$$\hat{r}_i(s) = \max \left(\left\{ \mathbb{I}_k(G_i^{3D}) \cdot \frac{f_k \cdot s_k}{d_k} \right\}_{k=1}^K \right), \quad (6)$$

where s_k is the scale factor of the k -th camera. ρ is folded into hyperparameter γ when constraining the highest signal frequency using Eq. 3. By integrating the scale factor, our method ensures a more adaptive and accurate Gaussian bandwidth constraint across varying resolutions.

2D Scale-Aware Mip Filter. When shading pixels with the Gaussian signal, the integration window size depends on the number of pixels (see Fig. 2(a)), with more pixels implying a smaller window. To study the effect of window size

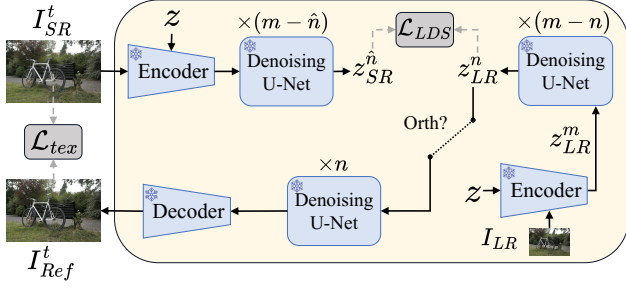


Figure 3: Generative prior-guided optimization, where generative priors are leveraged to constrain details in rendered HR views. To alleviate view inconsistency introduced by generative priors, optimization is conducted in the latent space, and texture supervision is applied only to orthogonal views.

on integration, we analyze the 1D approximation error by comparing the CDF of a standard Gaussian with that of the vanilla 2D Mip filter. As shown by the blue curve in Fig. 2(b), as the actual window size increases, integrating with a fixed window area accumulates approximation error. This reveals the limitation of static integration for scale-dependent signal capture. Therefore, we adaptively set the window size by choosing the 2D Gaussian filter variance as $\varepsilon_k = \frac{\varepsilon}{s_k}$. The scale-aware 2D Gaussian filter is then incorporated into 2D Gaussian signals with:

$$G_i^{2D}(\hat{\mathbf{x}})_{mip} = \sqrt{\frac{|\hat{\Sigma}_i|}{|\hat{\Sigma}_i + \varepsilon_k \cdot \mathbf{I}|}} e^{-\frac{1}{2}(\hat{\mathbf{x}} - \hat{\mu}_i)^\top (\hat{\Sigma}_i + \varepsilon_k \cdot \mathbf{I})^{-1} (\hat{\mathbf{x}} - \hat{\mu}_i)}. \quad (7)$$

The orange curve in Fig. 2(b) shows that the proposed scale-aware Mip filter consistently keeps a lower approximation error compared to the vanilla Mip filter. This highlights its effectiveness in precisely capturing signals (see Sec. 5.5).

4.2 Generative Prior-guided Optimization

For Arbi-3DGSR, only input LR views are available. To constrain fine details in the rendered HR views, we use generative priors from StableSR (Wang et al. 2024) as texture-rich references. However, directly optimizing with generated reference causes view inconsistencies due to the diverse and stochastic nature of the generation task (see Sec. 5.5). To this end, we propose the generative latent distillation and optimize with orthogonal reference refinement, ensuring both high perceptual quality and structural coherence across views.

Generative Latent Distillation. In Fig. 3, given an LR view I_{LR} , we sample a random Gaussian noise z , and embed it into the latent space conditioned on I_{LR} . This results in a noisy latent z_{LR}^m , where m is the initial denoising timestep. z_{LR}^m is then iteratively denoised through a UNet-based diffusion process $z_{LR}^{m-1} = UNet(z_{LR}^m, m)$, progressively producing an intermediate latent z_{LR}^n that is enriched with structural information. Meanwhile, the scale-aware rendered view I_{SR}^t at t -th training stage undergoes the same conditional diffusion process. This yields an intermediate latent $z_{SR}^{\hat{n}}$, where \hat{n} is an early denoising timestep ($\hat{n} > n$). z_{LR}^n therefore contains richer structural information from generative priors. To inject structural information from z_{LR}^n into $z_{SR}^{\hat{n}}$ and guide the optimization, we introduce the Latent Distillation Sampling




Figure 4: Progressive super-resolving. The training process is divided into multiple stages, with each stage following the same mechanism while progressively rendering higher-resolution views. Structural loss is applied between adjacent stages to ensure consistency across scale factors.

(LDS) loss \mathcal{L}_{LDS} that computes the gradient as follows:

$$\nabla_{\theta} \mathcal{L}_{LDS}(\theta) = \mathbb{E}_{\hat{n}} \left[w(\hat{n}) \cdot (\epsilon_{\phi}(z_{SR}^{\hat{n}}; I_{SR}^t, \hat{n}) - \epsilon_{\phi}(z_{LR}^n; I_{LR}, n)) \frac{\partial I_{SR}^t}{\partial \theta} \right], \quad (8)$$

where $\epsilon_{\phi}(\cdot)$ indicates predicting noise with UNet. $w(\hat{n})$ is a timestep \hat{n} related weighting function. θ denotes learnable parameters of 3D Gaussian primitives. Unlike the Score Distillation Sampling loss (Poole et al. 2023) that compares predicted noise with attached Gaussian noise of the same timestep, Eq. 8 minimizes the noise discrepancy between asynchronous latents $z_{SR}^{\hat{n}}$ and z_{LR}^n . This encourages $z_{SR}^{\hat{n}}$ to approximate the detailed texture in z_{LR}^n , while fully leveraging the structural information of LR views. Furthermore, minimizing the noise discrepancy instead of pixel discrepancy effectively provides structural supervision while tolerating pixel misalignment introduced by generative priors. We provide further analysis and theoretical derivation of LDS loss in Sec. 5.5 and supplementary materials, respectively.

Orthogonal Reference Refinement. While generative latent distillation provides inconsistency-tolerant supervision, explicit pixel-wise supervision is crucial for enhancing fine-grained texture details. To address potential inconsistencies between adjacent views, we employ an orthogonal view strategy: for each scene, we identify a subset of views that are approximately orthogonal (*i.e.*, the number of views is scene-dependent) for pixel-wise refinement. As shown in Fig. 3, for orthogonal views, the latent variable z_{LR}^n is further denoised and decoded to obtain an HR reference I_{Ref}^t . The texture loss \mathcal{L}_{tex} is conducted with Eq. 9 below:

$$\mathcal{L}_{tex} = \mathbb{I}_{ortho} \cdot \|I_{SR}^t - I_{Ref}^t\|^2, \quad (9)$$

where \mathbb{I}_{ortho} indicates whether current view is orthogonal. The orthogonal reference refinement encourages learning from non-overlapping views, avoiding conflicting information and preserving geometric consistency.

4.3 Progressive Super-Resolving

As we use a single 3D model for arbitrary-scale HR rendering, each output must remain structurally consistent with LR views. Otherwise, even slight misalignments magnify into visible warping or registration errors. To this end, training is divided into multiple stages to progressively accommodate various scale factors (see Fig. 4). Each stage adopts the same

Method	×2			×4			×8			×3.5			×5.7			
	PSNR↑	SSIM↑	FID↓	PSNR↑	SSIM↑	FID↓	PSNR↑	SSIM↑	FID↓	PSNR↑	SSIM↑	FID↓	PSNR↑	SSIM↑	FID↓	
Blender	3DGS	19.66	0.834	129.567	17.84	0.789	208.166	16.86	0.797	244.212	18.10	0.792	198.147	17.28	0.790	229.619
	Bicubic	20.36	0.853	132.513	19.67	0.831	178.170	19.24	0.831	182.420	19.76	0.833	176.546	19.46	0.829	181.065
	ArbiSR	19.34	0.827	152.266	18.23	0.803	185.134	17.98	0.817	174.873	18.33	0.803	185.430	18.08	0.808	179.436
	StableSR	18.63	0.788	120.385	18.25	0.785	95.086	17.98	0.797	94.957	18.31	0.785	95.720	18.12	0.790	94.853
	Mip-Splatting	23.33	0.896	78.907	22.25	0.870	109.435	21.57	0.858	113.124	22.40	0.874	108.101	21.92	0.862	111.875
	Analytic-Splatting	25.23	0.915	103.452	23.57	0.873	141.302	22.51	0.847	146.244	23.82	0.881	139.486	23.02	0.857	144.521
GaussianSR	23.25	0.885	94.017	23.03	0.868	118.021	22.37	0.856	117.618	23.13	0.871	117.458	22.73	0.861	118.465	
	Ours	25.60	0.925	66.515	24.32	0.899	86.270	23.40	0.879	87.984	24.53	0.904	85.788	23.87	0.888	87.429
Mip-NeRF360	3DGS	23.55	0.710	49.057	20.98	0.610	68.864	19.92	0.632	85.628	21.32	0.619	64.968	20.33	0.609	77.907
	Bicubic	25.92	0.738	44.731	25.09	0.687	47.441	24.86	0.730	47.551	25.20	0.689	46.805	24.96	0.699	47.450
	ArbiSR	25.14	0.727	56.838	23.92	0.671	57.471	23.87	0.714	54.983	24.02	0.674	57.714	23.86	0.683	56.061
	StableSR	20.85	0.480	107.550	20.54	0.507	106.957	20.45	0.605	106.950	20.58	0.495	107.272	20.49	0.551	106.968
	Mip-Splatting	26.10	0.750	35.931	25.17	0.696	38.950	24.51	0.718	41.025	25.30	0.700	38.294	25.02	0.702	39.089
	Analytic-Splatting	25.80	0.754	27.536	24.02	0.652	28.421	23.04	0.615	28.500	24.27	0.668	28.180	23.41	0.621	28.734
GaussianSR	24.93	0.660	103.160	24.33	0.650	105.013	24.10	0.710	104.943	24.42	0.646	104.548	24.20	0.674	105.031	
	Ours	26.23	0.764	36.524	25.18	0.703	38.526	24.85	0.725	38.645	25.32	0.709	37.924	24.99	0.704	38.606
Tanks&Temples	3DGS	19.45	0.711	72.678	16.24	0.516	135.005	14.86	0.468	176.733	16.67	0.541	121.909	15.43	0.479	161.616
	Bicubic	21.80	0.742	79.308	20.19	0.599	127.444	19.27	0.554	141.343	20.45	0.619	120.743	19.74	0.567	137.071
	ArbiSR	20.85	0.719	148.429	18.90	0.581	188.368	18.40	0.547	165.008	19.11	0.598	188.735	18.66	0.557	176.148
	StableSR	19.41	0.653	90.379	18.54	0.538	86.072	17.82	0.498	85.640	18.69	0.553	86.732	18.19	0.510	84.504
	Mip-Splatting	22.78	0.806	53.183	20.97	0.667	86.023	19.90	0.592	99.700	21.27	0.690	79.479	20.44	0.621	95.086
	Analytic-Splatting	21.75	0.754	73.966	19.42	0.577	127.344	18.20	0.498	143.717	19.78	0.605	118.811	18.75	0.525	137.808
GaussianSR	21.90	0.738	80.779	20.63	0.623	102.969	19.60	0.560	112.046	20.91	0.643	97.359	20.13	0.584	109.541	
	Ours	22.94	0.821	48.659	21.14	0.686	75.591	19.99	0.600	89.928	21.47	0.710	69.519	20.57	0.635	85.057
Deep Blending	3DGS	25.57	0.794	159.958	23.70	0.736	222.691	22.81	0.748	258.275	23.99	0.741	211.857	23.19	0.738	243.801
	Bicubic	26.40	0.817	140.654	25.60	0.793	173.301	25.27	0.814	185.594	25.75	0.793	169.449	25.46	0.802	180.383
	ArbiSR	26.09	0.817	161.001	24.87	0.787	195.876	24.62	0.808	200.379	25.03	0.787	192.530	24.74	0.797	199.635
	StableSR	24.31	0.755	131.813	23.81	0.729	127.640	23.57	0.749	128.898	23.90	0.729	129.270	23.70	0.736	128.558
	Mip-Splatting	26.75	0.830	137.672	25.90	0.805	174.613	25.52	0.818	189.308	26.06	0.806	168.978	25.73	0.810	182.643
	Analytic-Splatting	26.26	0.805	159.607	25.20	0.755	201.319	24.70	0.757	217.368	25.39	0.761	195.165	24.95	0.752	210.790
GaussianSR	26.89	0.843	134.990	26.30	0.826	163.499	25.91	0.832	172.887	26.45	0.827	159.877	26.13	0.829	169.988	
	Ours	27.44	0.861	106.533	26.57	0.835	140.633	26.14	0.836	154.810	26.75	0.837	134.836	26.38	0.835	148.951

Table 1: Quantitative results on Blender, Mip-NeRF360, Tanks&Temples, and Deep Blending dataset, with the best and second-best results highlighted in red and yellow, respectively. The proposed method achieves the overall best performance, demonstrating its effectiveness in rendering high-fidelity results and preserving consistency across different scale factors.

working mechanism and loss function, with the maximum scale factor increasing at each stage, enabling the model to gradually refine details at higher resolutions. At the t -th stage, the 3D Gaussian primitives are initialized from those learned in the previous $(t-1)$ -th stage. For each HR view, the rendering scale factor is randomly selected from $\{s^1, s^2, \dots, s^t\}$, where s^t is the largest scale factor used at the t -th stage. To maintain structural consistency across varying scale factors, a structure loss is imposed between current t -th stage HR view and corresponding lower-resolution view from previous $(t-1)$ -th stage. Specifically, given the current scale factor s^i , the previous stage renders with s^{i-1} to ensure a smooth transition between scales. The structure loss \mathcal{L}_{str} is conducted between the rendered outputs I_{SR}^t and I_{SR}^{t-1} as follows:

$$\mathcal{L}_{str} = (1 - \lambda) \mathcal{L}_{MSE} \left(\mathcal{D} \left(I_{SR}^t, \frac{s^i}{s^{i-1}} \right), I_{SR}^{t-1} \right) + \lambda \mathcal{L}_{D-SSIM} \left(\mathcal{D} \left(I_{SR}^t, \frac{s^i}{s^{i-1}} \right), I_{SR}^{t-1} \right) \quad (10)$$

where $\mathcal{D}(I_{SR}^t, \frac{s^i}{s^{i-1}})$ indicates downsampling the HR I_{SR}^t

by $\frac{s^i}{s^{i-1}}$. λ is the hyperparameter controlling the balance between each loss term. Applying structural loss between adjacent stages helps preserve structural similarity, ensuring smooth transitions and continuity across different scales. As a result, our overall objective function is described below:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{LDS} + \lambda_2 \mathcal{L}_{tex} + \lambda_3 \mathcal{L}_{str}, \quad (11)$$

where λ_1 , λ_2 and λ_3 denote weights of each loss term.

5 Experiments

5.1 Experimental Settings

Datasets. Experiments are conducted on four benchmark datasets: Blender (Mildenhall et al. 2021), Mip-NeRF360 (Barron et al. 2022), Tanks&Temples (Knapitsch et al. 2017), and Deep Blending (Hedman et al. 2018), following the same train-test splits as 3DGS (Kerbl et al. 2023). LR input views are generated by downsampling the original full-resolution images by a factor of 8 using Bicubic interpolation. $\times 2$, $\times 4$ and $\times 8$ are adopted as the upscale factors across three training stages. Notably, original full-resolution images

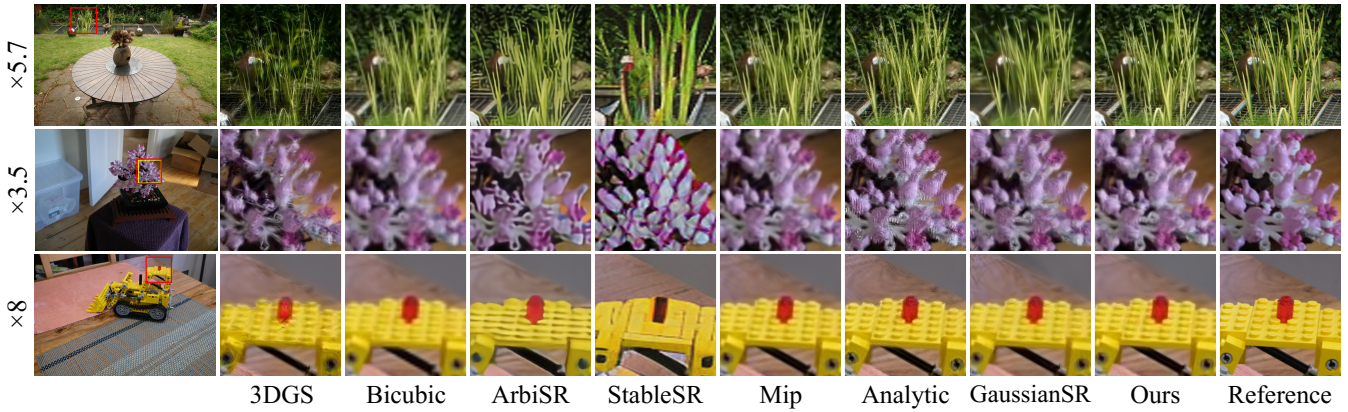


Figure 5: Qualitative comparisons on the Mip-NeRF360 dataset, where Mip and Analytic denote Mip-Splatting and Analytic-Splatting, respectively. Please zoom in for better results. As can be seen, 3DGS contains aliasing artifacts (e.g., 1st column). StableSR changes the contents of the rendered view (e.g., lego in the 3rd row). Analytic-Splatting and GaussianSR generate high-frequency artifacts. In contrast, the proposed method effectively renders high-fidelity results with rich details.

are not used during training, as ground-truth is unavailable in the super-resolution scenario. To assess generalization ability across arbitrary scales, we evaluate both integer ($\times 2$, $\times 4$ and $\times 8$) and non-integer ($\times 3.5$ and $\times 5.7$) scale factors, with corresponding ground truth views obtained by downsampling the original HR images to the target resolution.

Metrics. PSNR and SSIM are employed as error-based metrics to evaluate the performance. FID (Heusel et al. 2017) is adopted to assess the perceptual quality. The original full-resolution views serve as the reference distribution.

Baselines. We conduct comparisons with vanilla 3DGS (Kerbl et al. 2023), cascaded baselines, anti-aliasing techniques (i.e., Mip-Splatting (Yu et al. 2024b) and Analytic-Splatting (Liang et al. 2025)), and 3D super-resolution GaussianSR (Yu et al. 2024a). Vanilla 3DGS and anti-aliasing baselines are trained on LR views and conduct HRNVS of target scale factors. For cascaded solutions, LR novel views are first rendered with vanilla 3DGS and then upsampled with upsamplers, including Bicubic interpolation, arbitrary-scale ArbiSR (Li et al. 2025b), and generative StableSR (Wang et al. 2024). For simplicity, cascaded methods are denoted by corresponding upsamplers (e.g., Bicubic). We convert GaussianSR into an arbitrary-scale solution by training with random upscale factors.

5.2 Quantitative Results

Quantitative results are reported in Tab. 1. Our method shows an overall superiority over baselines across four benchmarks, demonstrating effectiveness in rendering realistic HR outputs. Specifically, vanilla 3DGS suffers from aliasing artifacts and obtains the worst performance. While Bicubic achieves favorable PSNR/SSIM, it lacks details and shows poor performance on FID. ArbiSR and StableSR provide limited improvement on PSNR/SSIM due to altered contents (see Fig. 5). Our method delivers superior perceptual quality over Mip-Splatting (by FID), and surpasses Analytic-Splatting in reconstructing high-fidelity results (e.g., a PSNR gain of 1.81 dB on Mip-NeRF360 at $\times 8$). Compared to GaussianSR, our method is consistently superior over all metrics.

Method	Rendering (ms)	Throughput (FPS)	Storage Size (GB)	Training (min)	Memory (MB)
3DGS	13	74	0.99	10	718
Bicubic	34	29	-	-	-
ArbiSR	3225	0.31	-	-	-
StableSR	10890	0.13	-	-	-
Mip	19	52	0.99	12	858
Analytic	33	30	1.20	29	822
GaussianSR	8	126	0.56	256	8274
Ours	12	85	0.79	57	7160

Table 2: Efficiency analysis measured on a single NVIDIA A6000 GPU. Rendering time and throughput are evaluated at an output resolution of 1920×1080 . Mip and Analytic denote Mip-Splatting and Analytic-Splatting, respectively.

5.3 Qualitative Results

We provide qualitative comparisons on Mip-NeRF360 in Fig. 5, where full-resolution images are included as reference. As can be seen, 3DGS suffers from aliasing artifacts, leading to severely degraded structures and erosion effects (e.g., the grass in the 1st row). Bicubic and Mip-Splatting produce over-smoothed outputs (e.g., the bonsai in the 3rd row). ArbiSR introduces distorted textures that compromise the visual quality. Despite StableSR providing visually appealing results, it alters the content and deviates from the original LR input views. Analytic-Splatting and GaussianSR hallucinate high-frequency artifacts that resemble textures. In contrast, our method consistently delivers high-quality results, while faithfully preserving the original structures. We include more qualitative comparisons in the supplementary materials.

5.4 Efficiency Analysis

Efficiency comparisons on rendering time, throughput, storage cost, training time, and GPU memory usage are included in Tab. 2. Our method introduces no extra computational overhead during rendering, achieving a throughput of 85 FPS for 1080p rendering, significantly faster than cascaded solutions (e.g., $908\times$ faster than StableSR). With fewer 3D Gaussians (as reflected by reduced storage size), GaussianSR

Mip-NeRF360	×2			×4			×8			×3.5			×5.7		
	PSNR↑	SSIM↑	FID↓	PSNR↑	SSIM↑	FID↓	PSNR↑	SSIM↑	FID↓	PSNR↑	SSIM↑	FID↓	PSNR↑	SSIM↑	FID↓
w/o 3D-SASF	26.13	0.757	41.576	24.85	0.685	43.804	24.39	0.697	43.888	25.04	0.693	43.149	24.57	0.681	43.881
w/o 2D-SAMF	25.53	0.736	36.858	24.83	0.679	39.613	24.61	0.703	39.774	24.93	0.684	39.006	24.71	0.682	39.610
w/o scale-aware	26.10	0.758	37.052	25.17	0.698	39.525	24.87	0.721	39.635	25.30	0.703	38.901	25.00	0.700	39.523
w/o PSR	26.03	0.753	37.921	24.51	0.663	39.748	23.91	0.664	39.873	24.74	0.674	39.124	24.14	0.651	39.797
w/o GPO	25.23	0.667	99.685	24.51	0.656	101.618	24.27	0.714	101.451	24.61	0.652	101.204	24.37	0.678	101.530
w/o Orth	25.97	0.764	39.392	25.11	0.713	40.688	24.85	0.741	40.573	25.24	0.717	40.106	24.96	0.718	40.633
Pseudo HR	23.96	0.593	111.149	23.36	0.600	111.574	23.19	0.683	111.215	23.44	0.592	111.584	23.26	0.636	111.526
SDS loss	23.52	0.698	72.640	22.91	0.655	75.134	22.71	0.697	75.830	22.99	0.657	74.753	22.79	0.667	75.336
Ours	26.23	0.764	36.524	25.18	0.703	38.526	24.85	0.725	38.645	25.32	0.709	37.924	24.99	0.704	38.606

Table 3: Ablation studies on Mip-NeRF360 dataset, where the best and second best results are highlighted with red and yellow, respectively. 3D-SASF, 2D-SAMF, PSR, GPO, Orth denote the 3D scale-aware smoothing filter, 2D scale-aware Mip filter, progressive super-resolving, generative prior-guided optimization and orthogonal reference refinement, respectively.



Figure 6: Qualitative comparisons with typical optimization strategies at the scale factor of $\times 4$. Replacing our generative prior-guided optimization with pseudo HR supervision and SDS loss results in blurriness and color distortion.

and our method naturally render faster over Mip-Splatting and Analytic-Splatting. Notably, our method still delivers superior performance in Tab. 1 (e.g., a PSNR gain of 2.13 dB over Mip-Splatting on Blender at $\times 3.5$), highlighting its effectiveness in learning compact and expressive 3D scene representations. During training, generative prior-guided optimization inevitably increases the training time and GPU memory usage. However, the overhead remains affordable on commonly used GPUs. For instance, training takes only 57 min per scene with approximately 7 GB of memory usage.

5.5 Ablation Studies

We assess each component’s contribution by removing it. Then we examine the effectiveness of scale-awareness in the 2D scale-aware Mip filter. We further compare our generative prior-guided optimization with typical optimization strategies. Ablation studies are conducted on both Mip-NeRF360 and Tanks&Temples (see supplementary materials) for generality.

3D Scale-Aware Smooth Filter. As shown in Tab. 3, removing 3D scale-aware smooth filter (denoted as w/o 3D-SASF) results in degraded performance. For instance, it leads to a PSNR drop of 0.46 dB on Mip-NeRF360 at $\times 8$.

2D Scale-Aware Mip Filter. As shown in Tab. 3, removing 2D scale-aware Mip filter (denoted as w/o 2D-SAMF) leads to a notable performance drop. For instance, the PSNR shows a drop of 0.7 dB at $\times 2$. Unlike the vanilla 2D Mip filter (Yu et al. 2024b), our proposed filter adaptively adjusts the integration window based on the scale factor, enabling more accurate pixel shading (as outlined Sec.4.1). To investigate the effectiveness of scale-awareness, we replace 2D

scale-aware Mip filter with the fixed-window 2D Mip filter (denoted as w/o SA). As shown in Tab. 3, this results in sub-optimal performance, especially in realism-based FID. Visual comparisons are included in the supplementary materials.

Progressive Super-Resolving. To demonstrate the effectiveness of progressive super-resolving, we replace it with a mix-training strategy of random scale factors (denoted as w/o PSR). As shown in the 4th row of Tab. 3, training without progressive super-resolving leads to suboptimal results.

Generative Prior-Guided Optimization. As shown in Tab. 3, removing generative prior-guided optimization (w/o GPO) leads to a significant performance drop (e.g., a PSNR drop of 1 dB at $\times 2$). Without generative priors, outputs exhibit reduced realism, as evidenced by increased FID scores.

Comparison of Typical Optimization Strategies. We further compare our GPO with two widely used optimization strategies: (1) *pseudo HR supervision*, which imposes pixel-wise loss with pseudo HR views generated by pretrained SR models (Xie et al. 2024; Feng et al. 2024); and (2) *Score Distillation Sampling (SDS) loss*, which leverages generative priors via a denoising process (Poole et al. 2023). As shown in Tab. 3, pseudo HR supervision yields suboptimal results, particularly in perceptual realism (as reflected by FID), due to inconsistencies across views that hinder detail reconstruction and cause blur (see Fig. 6). SDS loss also shows notable drops in PSNR/SSIM, exhibiting color distortions in Fig. 6. In contrast, our GPO consistently delivers superior performance, effectively refining details and keeping fidelity.

6 Conclusion

This paper introduces an integrated framework for arbitrary-scale 3D Gaussian super-resolution. To address the key challenges, our framework incorporates three core components: scale-aware rendering, which enables adaptive HR rendering across varying scale factors; generative prior-guided optimization, which constrains fine details in the absence of ground-truth HR views; and progressive super-resolving, which gradually super-resolves 3D models to preserve content consistency across varying scale factors. Experimental results on four benchmarks demonstrate the effectiveness of our method in rendering high-quality arbitrary-scale HR views with a single 3D model. Ablation studies further verify the effectiveness of each component in our framework.

References

- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2022. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5470–5479.
- Bondarets, A.-A.; Rumezhak, T.; Plakhtii, H.; Muliak, A.; and Karpiv, V. 2024. 3D-GSR: 3D Super-Resolution with Multi-View Adapter and Gaussian Splatting. In *ECCV 2024 Workshop on Wild 3D: 3D Modeling, Reconstruction, and Generation in the Wild*.
- Chen, Z.; Chen, Y.; Liu, J.; Xu, X.; Goel, V.; Wang, Z.; Shi, H.; and Wang, X. 2022. VideoInr: Learning video implicit neural representation for continuous space-time super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2047–2057.
- Feng, X.; He, Y.; Wang, Y.; Yang, Y.; Li, W.; Chen, Y.; Kuang, Z.; Fan, J.; Jun, Y.; et al. 2024. Srgs: Super-resolution 3d gaussian splatting. *arXiv preprint arXiv:2404.10318*.
- Fu, H.; Peng, F.; Li, X.; Li, Y.; Wang, X.; and Ma, H. 2024. Continuous Optical Zooming: A Benchmark for Arbitrary-Scale Image Super-Resolution in Real World. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3035–3044.
- Fu, Y.; Li, Y.; Li, C.; Saragih, J.; Zhang, P.; Dai, X.; and Lin, Y. C. 2023. Auto-CARD: Efficient and Robust Codec Avatar Driving for Real-time Mobile Telepresence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21036–21045.
- Grant, G.; Burton, R.; Grafton, E.; Della-Bosca, D.; Ditcham, R.; and Humphreys, L. 2023. Meta-patients: Using mixed reality patients and an ai framework for simulating life-like clinical examinations. In *Augmented Reality and Artificial Intelligence: The Fusion of Advanced Technologies*, 193–210. Springer.
- Han, Y.; Yu, T.; Yu, X.; Xu, D.; Zheng, B.; Dai, Z.; Yang, C.; Wang, Y.; and Dai, Q. 2024. Super-NeRF: View-consistent Detail Generation for NeRF Super-resolution. *IEEE Transactions on Visualization and Computer Graphics*, 1–14.
- Hedman, P.; Philip, J.; Price, T.; Frahm, J.-M.; Drettakis, G.; and Brostow, G. 2018. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (ToG)*, 37(6): 1–15.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Hu, J.; Xia, B.; Chen, B.; Yang, W.; and Zhang, L. 2024. GaussianSR: High Fidelity 2D Gaussian Splatting for Arbitrary-Scale Image Super-Resolution. *arXiv preprint arXiv:2407.18046*.
- Hu, X.; Mu, H.; Zhang, X.; Wang, Z.; Tan, T.; and Sun, J. 2019. Meta-SR: A magnification-arbitrary network for super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1575–1584.
- Huang, D.-J.; Chou, Z.-T.; Wang, Y.-C. F.; and Sun, C. 2024. ASSR-NeRF: Arbitrary-Scale Super-Resolution on Voxel Grid for High-Quality Radiance Fields Reconstruction. *arXiv preprint arXiv:2406.20066*.
- Huang, X.; Li, W.; Hu, J.; Chen, H.; and Wang, Y. 2023. RefSR-NeRF: Towards High Fidelity and Super Resolution View Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8244–8253.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4): 139–1.
- Knapitsch, A.; Park, J.; Zhou, Q.-Y.; and Koltun, V. 2017. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4): 1–13.
- Ko, H.-k.; Park, D.; Park, Y.; Lee, B.; Han, J.; and Park, E. 2024. Sequence Matters: Harnessing Video Models in Super-Resolution. *arXiv preprint arXiv:2412.11525*.
- Lee, J. L.; Li, C.; and Lee, G. H. 2024. DiSR-NeRF: Diffusion-Guided View-Consistent Super-Resolution NeRF. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20561–20570.
- Li, B.; Li, X.; Liang, Y.; Wang, S.; and Song, K.-B. 2025a. Leveraging Latent Diffusion in 3D Gaussian Splatting for Novel View Synthesis. In *International Conference on Multimedia Modeling*, 150–157. Springer.
- Li, Z.; Li, M.; Fan, J.; Chen, L.; Tang, Y.; Lu, J.; and Zhou, J. 2025b. Learning Dual-Level Deformable Implicit Representation for Real-World Scale Arbitrary Super-Resolution. In *European Conference on Computer Vision*, 352–368. Springer.
- Li, Z.; Liu, H.; Shang, F.; Liu, Y.; Wan, L.; and Feng, W. 2024. SAVSR: Arbitrary-Scale Video Super-Resolution via a Learned Scale-Adaptive Network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 3288–3296.
- Liang, Z.; Zhang, Q.; Hu, W.; Zhu, L.; Feng, Y.; and Jia, K. 2025. Analytic-splatting: Anti-aliased 3d gaussian splatting via analytic integration. In *European conference on computer vision*, 281–297. Springer.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Nyquist, H. 1928. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2): 617–644.
- Poole, B.; Jain, A.; Barron, J. T.; and Mildenhall, B. 2023. DreamFusion: Text-to-3D using 2D Diffusion. In *The Eleventh International Conference on Learning Representations*.
- Shang, W.; Ren, D.; Zhang, W.; Fang, Y.; Zuo, W.; and Ma, K. 2025. Arbitrary-Scale Video Super-Resolution with Structural and Textural Priors. In *European Conference on Computer Vision*, 73–90. Springer.

Shannon, C. E. 1949. Communication in the presence of noise. *Proceedings of the IRE*, 37(1): 10–21.

Shen, Y.; Ceylan, D.; Guerrero, P.; Xu, Z.; Mitra, N. J.; Wang, S.; and Frühstück, A. 2025. Supergaussian: Repurposing video models for 3d super resolution. In *European Conference on Computer Vision*, 215–233. Springer.

Solmaz, S.; and Van Gerven, T. 2020. Integration of interactive CFD simulations with AR and VR for educational Use in CRE. In *Computer Aided Chemical Engineering*, volume 48, 2011–2016. Elsevier.

Song, X.; Zheng, J.; Yuan, S.; Gao, H.-a.; Zhao, J.; He, X.; Gu, W.; and Zhao, H. 2024. Sa-gs: Scale-adaptive gaussian splatting for training-free anti-aliasing. *arXiv preprint arXiv:2403.19615*.

Wang, C.; Wu, X.; Guo, Y.-C.; Zhang, S.-H.; Tai, Y.-W.; and Hu, S.-M. 2022. Nerf-sr: High quality neural radiance fields using supersampling. In *Proceedings of the 30th ACM International Conference on Multimedia*, 6445–6454.

Wang, J.; Yue, Z.; Zhou, S.; Chan, K. C.; and Loy, C. C. 2024. Exploiting diffusion prior for real-world image super-resolution. *International Journal of Computer Vision*, 1–21.

Weng, S.-Y.; Yuan, H.; Xu, Y.-S.; Huang, C.-C.; and Chiu, W.-C. 2024. Best of Both Worlds: Learning Arbitrary-scale Blind Super-Resolution via Dual Degradation Representations and Cycle-Consistency. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1547–1556.

Xie, S.; Wang, Z.; Wang, X.; Zhu, Y.; Pan, C.; and Dong, X. 2024. SuperGS: Super-Resolution 3D Gaussian Splatting Enhanced by Variational Residual Features and Uncertainty-Augmented Learning. *arXiv:2410.02571*.

Xu, R.; Yao, M.; Li, Y.; Zhang, Y.; and Xiong, Z. 2025. High-Resolution and Few-Shot View Synthesis from Asymmetric Dual-Lens Inputs. In *European Conference on Computer Vision*, 218–235. Springer.

Yan, Z.; Low, W. F.; Chen, Y.; and Lee, G. H. 2024. Multi-scale 3d gaussian splatting for anti-aliased rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20923–20931.

Yoon, Y.; and Yoon, K.-J. 2023. Cross-Guided Optimization of Radiance Fields With Multi-View Image Super-Resolution for High-Resolution Novel View Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12428–12438.

Yu, X.; Zhu, H.; He, T.; and Chen, Z. 2024a. GaussianSR: 3D Gaussian Super-Resolution with 2D Diffusion Priors. *arXiv preprint arXiv:2406.10111*.

Yu, Z.; Chen, A.; Huang, B.; Sattler, T.; and Geiger, A. 2024b. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19447–19456.

Zhao, Y.; Teng, Q.; Chen, H.; Zhang, S.; He, X.; Li, Y.; and Sheriff, R. E. 2024. Activating More Information in Arbitrary-Scale Image Super-Resolution. *IEEE Transactions on Multimedia*, 26: 7946–7961.