

# LiDAR-GS++: Improving LiDAR Gaussian Reconstruction via Diffusion Priors

Qifeng Chen<sup>1</sup>, Jiarun Liu<sup>1</sup>, Rengan Xie<sup>2</sup>, Tao Tang<sup>3</sup>, Sicong Du<sup>1</sup>, Yiru Zhao<sup>1</sup>,  
Yuchi Huo<sup>2</sup>, Sheng Yang<sup>1\*</sup>

<sup>1</sup>Unmanned Vehicle Dept., CaiNiao Inc., Alibaba Group

<sup>2</sup>State Key Laboratory of CAD&CG Zhejiang University

<sup>3</sup>Sun Yat-sen University

{cqf7419, shengyang93fs}@gmail.com

## Abstract

Recent GS-based rendering has made significant progress for LiDAR, surpassing Neural Radiance Fields (NeRF) in both quality and speed. However, these methods exhibit artifacts in extrapolated novel view synthesis due to the incomplete reconstruction from single traversal scans. To address this limitation, we present LiDAR-GS++, a LiDAR Gaussian Splatting reconstruction method enhanced by *diffusion priors* for real-time and high-fidelity re-simulation on public urban roads. Specifically, we introduce a controllable LiDAR generation model conditioned on coarsely extrapolated rendering to produce extra geometry-consistent scans and employ an effective distillation mechanism for expansive reconstruction. By extending reconstruction to under-fitted regions, our approach ensures global geometric consistency for extrapolative novel views while preserving detailed scene surfaces captured by sensors. Experiments on multiple public datasets demonstrate that LiDAR-GS++ achieves state-of-the-art performance for both interpolated and extrapolated viewpoints, surpassing existing GS and NeRF-based methods.

**Code** — [https://github.com/CN-ADLab/LiDAR\\_GS\\_plus](https://github.com/CN-ADLab/LiDAR_GS_plus)

## Introduction

Recent advances in closed-loop simulation (Hu et al. 2023a; Zeng et al. 2025; Ren et al. 2025) are continuously improving infrastructures for developing end-to-end autonomous driving algorithms (Hu et al. 2023b; Chen et al. 2024a; Fu et al. 2025), where Gaussian Splatting (GS) (Kerbl et al. 2023) has emerged as a primary representation for reconstruction-based simulators, enabling real-time rendering and ensuring physical realism. Facilitated by advancements in GS methods for camera (Yan et al. 2024; Zhou et al. 2024) and LiDAR (Jiang et al. 2025; Chen et al. 2024b), developers can reconstruct urban road scenes using crowd-sourced driving clips, which are typically collected to improve casually encountered corner cases, thereby meeting the requirements of arbitrary view synthesis and creating a virtual regression testing.

While reconstruction-based simulators offer high realism and reusability, they are constrained by the original viewpoint distribution of the driving clip, and thus exhibit limited




Figure 1: LiDAR re-simulation methods, such as GS-LiDAR (Jiang et al. 2025) and LiDAR-GS (Chen et al. 2024b), encounter performance drops when rendering extrapolated views, e.g., lateral viewpoint shifting. In contrast, LiDAR-GS++ maintains stable and reliable performance by diffusing priors.

performance for trajectory *extrapolation* when compared to *interpolation*. While simulators for cameras (Wang et al. 2025; Fan et al. 2025; Ni et al. 2025a) have made notable progress in addressing this problem, recent works on LiDAR re-simulation have not yet explicitly addressed the quality issue of extrapolation, which typically degrades simulation quality in viewpoint-changing scenarios, such as lane-changing for an obstacle avoidance. The imbalanced development of camera and LiDAR simulators decelerates the evolution of high-level (Level-4) driving agents equipped with multi-modal sensors (Liao et al. 2025; Cao et al. 2025).

Recent advances in general controllable generative models (Rombach et al. 2022; Blattmann et al. 2023; Zhang, Rao, and Agrawala 2023) open up the opportunity to augment LiDAR Gaussian reconstruction with generated data, contingent upon addressing two major challenges: (1) *Authenticity*. Relying on cross-modal prompts such as texts, maps, and bounding boxes can generate promising objectives, but their scene geometry exhibit a noticeable domain gap compared to real scans that is ineffective for novel view synthesis; and (2) *Consistency*. Blending generated scans with existing real scans may cause hallucinations and contradictions in specific regions, resulting in degraded harmoniousness for expansive reconstruction.

To address these challenges, we present LiDAR-GS++ and insist that LiDAR-to-LiDAR controllable generation serves as an appropriate strategy for scene expansion. Based on existing reconstruction-based simulators (Chen et al. 2024b; Jiang et al. 2025), we propose a novel model for

\*Corresponding author

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

conditioned LiDAR scan generation, prompted by coarsely rendered scans at extrapolated viewpoints, which improves geometry coherence and eliminates the need for auxiliary cross-modal prompts. Furthermore, given such generated scans for performing expansive reconstruction, we noticed that fully blending generated scans with existing real scans introduces confusion in regions which were already well-fitted. Thus, we introduce a depth distortion-aware distillation strategy accordingly to integrate these scans into the extrapolated reconstruction harmoniously.

According to our experiments on multiple public datasets, LiDAR-GS++ has achieved state-of-the-art re-simulation quality on both *extrapolation* and *interpolation*. In conclusion, LiDAR-GS++ makes the following key contributions:

- We propose a LiDAR reconstruction method that expansively reconstructs neural 2DGS fields via diffusion priors, achieving state-of-the-art quality compared to previous GS-based and NeRF-based re-simulation methods in both extrapolation and interpolation.
- For generating extrapolated LiDAR scans as extra supervision, we propose a controllable LiDAR generation model guided by coarsely extrapolated GS rendering.
- For expansive reconstruction, we propose a depth distortion aware distillation strategy to effectively integrate generated scans and harmoniously avoid hallucinations.

## Related Work

**Render-based LiDAR Reconstruction.** Recently, the rapid advancement of NeRF has driven the development of a series of LiDAR re-simulation methods. NFL (Huang et al. 2023) and LiDAR-NeRF (Tao et al. 2024) pioneered differentiable frameworks that model LiDAR beams as rays, rendering depth and intensity analogous to RGB images. DyNFL (Wu et al. 2024a) and LiDAR4D (Zheng et al. 2024) address dynamic reconstruction by incorporating temporal conditioning and instance-level dynamic decomposition, respectively. Mars (Wu et al. 2023) integrates this capability into a comprehensive simulation system.

Due to the limited rendering efficiency of NeRF, several concurrent previous works have proposed to refactor the scene representation into GS: GS-LiDAR (Jiang et al. 2025) extends LiDAR4D to 2DGS representations, significantly improving inference speed. LiDAR-GS (Chen et al. 2024b) proposes a 3DGS-based differentiable laser beam splatting and considers the effect of ray direction and distance on rendering. LiDAR-RT (Zhou et al. 2025) extends the Boundary Volume Hierarchy (BVH) of GS-RT (Moenne-Loccoz et al. 2024) to LiDAR and implements differentiable ray tracing using the OptiX framework (Parker et al. 2010).

Despite these advances, existing methods have not explicitly addressed the challenge of extrapolation. To this end, we are the first to integrate diffusion priors to expand lidar Gaussian reconstruction. Our approach builds upon the efficient Gaussian attribute prediction of LiDAR-GS (Chen et al. 2024b), while substituting its 3DGS representation with neural 2DGS for enhanced geometric fidelity. Comparative experiments demonstrate that this modification leads to more effective and efficient scene reconstruction.

**Diffusion-based LiDAR Generation.** General diffusion models (Ho, Jain, and Abbeel 2020; Rombach et al. 2022; Ren et al. 2024) have been demonstrated to be applicable to LiDAR generation tasks by leveraging structured representations such as 2D range view images or 3D voxel grids. Early works, including LiDARGen (Zyrianov, Zhu, and Wang 2022) and R2DM (Nakashima and Kurazume 2024), adopt the DDPM framework for unconditional LiDAR generation. Building on latent diffusion models (LDM) (Rombach et al. 2022), methods such as RangeLDM (Hu, Zhang, and Hu 2024), LiDM (Ran, Guizilini, and Wang 2024), and Copilot4D (Zhang et al. 2024) further improve generative quality by compressing LiDAR into a latent space using VAE (Kingma and Welling 2014) or VQ-VAE (van den Oord, Vinyals, and Kavukcuoglu 2017), enabling more efficient training and generation.

Meanwhile, the evolution of conditional controllable generation modules (Rombach et al. 2022; Zhang, Rao, and Agrawala 2023; Peng et al. 2024) has further catalyzed the development for controllable LiDAR generation. LidarDM (Zyrianov et al. 2024) employs traffic layouts to generate actors and scene. TYP (Pan et al. 2025) generates LiDAR Bird’s Eye View (BEV) using bounding boxes and ego view. Text2LiDAR (Wu et al. 2024b) pairs LiDAR scans with textual descriptions, pioneering the task of text-to-LiDAR. OLiDM (Yan et al. 2025) enhances object-aware LiDAR synthesis with category- and position-level prompts. UniScene (Li et al. 2025) leverages an occupancy-based intermediate representation to enable consistent generation of LiDAR and camera.

While these works successfully leverage cross-modal prompts to generate desired objectives, the quality and consistency of the results against a specific driving scan remain insufficient for launching an expansive reconstruction. Given prompts of the same modality, our LiDAR-to-LiDAR controllable generation applies appropriate diffusion and denoising conditioned on the incompletely reconstructed extrapolated rendering. Consequently, combining reconstruction and diffusion brings the opportunity to progressively enlarge the accessible LiDAR novel view synthesis range based on limited initial real scans.

## Method

As illustrated in Fig. 2, we first project the point cloud to the range view and reconstruct the scene using neural 2DGS field from a single-pass driving clip. Then, we render a set of extrapolated degraded LiDAR scans by manually shifting the viewpoints, which are used as a condition for a pre-trained controllable LiDAR generative model to produce geometrically consistent additional supervision for expansive reconstruction. During expansive reconstruction, as generated and real scans are fused, we employ a Depth Distortion-Aware Distillation (DDAD) strategy on generative data to specifically correct under-fitted regions and mitigate negative effects on convergence areas during reconstruction.

### Gaussian Splatting Scene Representation

**Range view Projection.** Following standard practices in LiDAR research for structurally representing distributions,

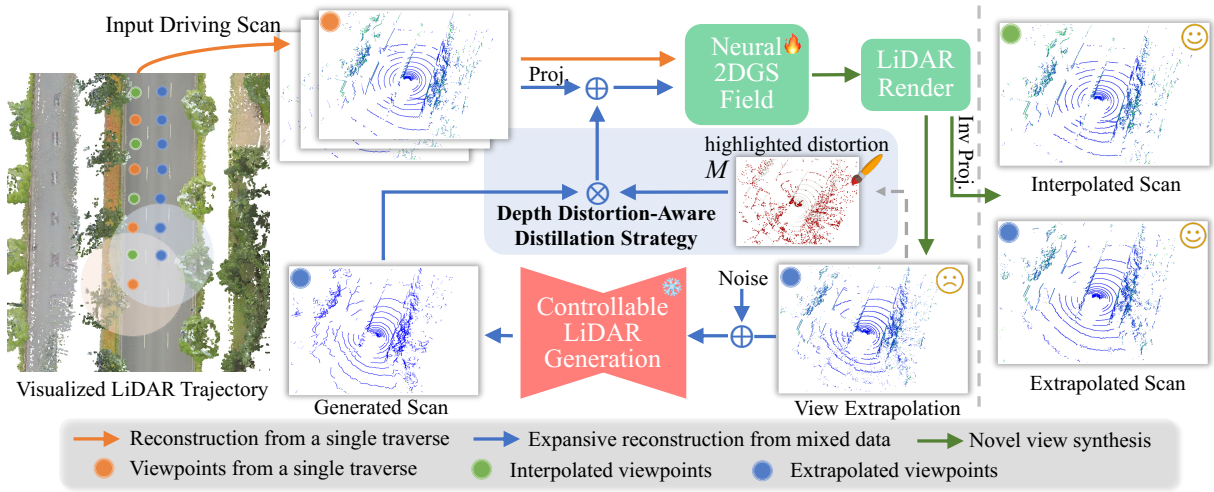


Figure 2: The workflow of LiDAR-GS++. Given an input driving scan, we first project the point cloud to the range view and reconstruct the scene using a neural 2DGS field from a single traverse. Then, we feed coarsely extrapolated rendering to a pre-trained LiDAR diffusion model to produce geometrically consistent extrapolated scans as extra supervision signals, and utilize a depth distortion-aware strategy to distill the GS representation for expansive reconstruction, where the distorted areas of the extrapolated view are highlighted in red. During inference, rendering from reconstructed Gaussians enables real-time, high-quality synthesis for both extrapolated and interpolated viewpoints. Subsequently, the range view can be converted into a point cloud format via inverse projection.

we convert LiDAR data to the range view, which comprises three channels: intensity  $\bar{\rho}$ , depth  $\bar{d}$  and ray-drop  $\bar{r}$  (Milioto et al. 2019; Kong et al. 2023; Chen et al. 2024b), where ray-drop  $\bar{r} \in \{0, 1\}$  indicate whether a beam has a return value. The projection process is as follows:

For each point  $(x, y, z)$  in the LiDAR coordinate, its corresponding pixel location  $(h, w)$  on the  $H \times W$  range view image, as:

$$\begin{pmatrix} h \\ w \end{pmatrix} = \begin{pmatrix} (1 - f(\phi, f_b)) \\ 0.5 \cdot (1 - \theta/\pi) \end{pmatrix} \cdot \begin{pmatrix} H \\ W \end{pmatrix}, \quad (1)$$

where  $\phi = \arcsin(z, \sqrt{x^2 + y^2})$  and  $\theta = \arctan(y, x)$  denote the vertical and horizontal angles, respectively. The function  $f(\phi, f_b)$  retrieves the position ratio closest to the vertical angle  $\phi$  from a pre-defined list of vertical angles  $f_b$  corresponding to each laser beam, as defined in the intrinsic parameters of the LiDAR. Inversely, re-visualizing the point cloud can be done simply by performing inverse projection:

$$(x, y, z) = d \cdot \mathbf{d} \triangleq d \cdot (\cos \theta \cos \phi, \sin \theta \cos \phi, \sin \phi), \quad (2)$$

where  $d$  stands for the distance-of-flight, and the ray-direction  $\mathbf{d}$  is determined by  $(\phi, \theta)$ .

**Neural 2DGS Field.** In light of the inherent distance- and direction-dependent attenuation of LiDAR returns (Hahner et al. 2021, 2022; Chen et al. 2024b) and the recent advances in geometric fidelity afforded by 2DGS (Huang et al. 2024), we propose a neural 2D Gaussian field for LiDAR scene modeling. It is composed of a set of 2D Gaussians  $\xi = \{\mathbf{x}, \mathbf{R}, \mathbf{S}, \rho, r, \alpha, \mathbf{v}_\xi\}$ , where center position  $\mathbf{x}$ , rotation quaternion  $\mathbf{R}$ , scale  $\mathbf{S}$ , intensity  $\rho$ , ray-drop probability  $r$ , opacity  $\alpha$  and feature token  $\mathbf{v}_\xi$  (32-dims initialized to zero),

respectively. Attributes of each Gaussian are predicted using four lightweight Multi-Layer Perceptrons (MLPs), receiving optimizable zero tokens  $\mathbf{v}_\xi$ , the local ray direction  $\mathbf{d}'$  and distance-of-flight  $d$  as input (see Fig. 3). This conditioning mechanism enables the network to predict direction- and distance-dependent attributes of 2D Gaussian. Additionally the Gaussian positions  $\mathbf{x}$  are refined by offsets, i.e.  $\mathbf{x} \leftarrow \mathbf{x} + \delta \mathbf{x}$ .




Figure 3: Design and workflow of the Neural 2DGS Field. The numbers on the vector represent the dimensions.

**LiDAR Rendering.** Given the free viewpoint and orientation for rendering, we follow the 2DGS (Huang et al. 2024) to compute the ray-splat intersections, and use range view rasterization to render the LiDAR’s depth, intensity, and ray-drop. Specifically, at first, each 2D Gaussian  $\xi$  in the world coordinate is parameterized in the local tangent plane:

$$\mathbf{P}(u, v) = \mathbf{W}\mathbf{H}(u, v, 1)^\top, \quad \mathbf{H} \triangleq \begin{bmatrix} s_u \mathbf{t}_u & s_v \mathbf{t}_v & \mathbf{x} \\ 0 & 0 & 1 \end{bmatrix} \in \mathbb{R}^{4 \times 3}, \quad (3)$$

where  $\mathbf{t}_u, \mathbf{t}_v \in \mathbb{R}^{1 \times 3}$  are two rotation components decomposed from rotation quaternion  $\mathbf{R}$ ,  $s_u$  and  $s_v$  from  $\mathbf{S} \triangleq \text{diag}(s_u, s_v)$  are their corresponding scaling components,

and  $\mathbf{W} \in \mathbb{R}^{4 \times 4}$  is the view projection matrix.  $\mathbf{P}(u, v)$  represents the homogeneous coordinates of a point on the local tangent plane in the world coordinate, while  $(u, v)$  denotes the relative position with respect to the 2D Gaussian center in Gaussian’s  $uv$  coordinate.

Then, we parameterize the ray of each pixel as the intersection of two planes,  $\mathbf{h}_u$  and  $\mathbf{h}_v$ , and the ray-splat intersection in Gaussian’s local tangent plane can be determined by solving:

$$\begin{aligned} [\mathbf{h}_u, \mathbf{h}_v]^\top \cdot \mathbf{P}(u, v) &= 0, \\ \mathbf{h}_u &= (\sin \phi, -\cos \phi, 0, 0)^\top, \\ \mathbf{h}_v &= (\sin \theta \cos \phi, \sin \theta \sin \phi, -\cos \theta, 0)^\top. \end{aligned} \quad (4)$$

Once the intersection  $(u, v)$  is determined, its probability distribution value can be calculated using the standard Gaussian function  $\mathbf{G}_\xi(u, v) = \exp(-(u^2 + v^2)/2)$ .

During the range view rasterization process, we record the Gaussian intersections that each pixel’s ray passes through, and use volume rendering integration to determine the values of the range view:

$$[\bar{\rho}, \bar{d}, \bar{r}] \leftarrow \sum_{i \in \mathbf{N}} [\rho_{\xi_i}, d_{\xi_i}, r_{\xi_i}] \alpha_{\xi_i} \mathbf{G}_{\xi_i} \prod_{j=1}^{i-1} (1 - \alpha_{\xi_j} \mathbf{G}_{\xi_j}), \quad (5)$$

where  $\bar{\rho}$ ,  $\bar{d}$ , and  $\bar{r}$  is the range view rendering of intensity, depth, and ray-drop, respectively.  $\mathbf{N}$  stands for the number of intersection on the ray.  $\mathbf{G}_{\xi_i}$  denotes probability distribution value of intersection on  $i$ -th Gaussian.

**Optimization.** We optimize the neural 2DGS field, including feature token  $\mathbf{v}_\xi$  of each Gaussian and parameters of four MLPs, by calculating loss between the input and rendered images:

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_d + \mathcal{L}_\rho + \mathcal{L}_r + \mathcal{L}_S, \\ \mathcal{L}_\rho &= (1 - \lambda_\rho) \cdot \mathcal{L}_1 + \lambda_\rho \cdot \mathcal{L}_{D\text{-SSIM}}, \end{aligned} \quad (6)$$

where the depth loss  $\mathcal{L}_d$  follows  $\mathcal{L}_1$  loss, and we set  $\lambda_\rho = 0.2$  for the intensity loss  $\mathcal{L}_\rho$ . The ray-drop probability loss  $\mathcal{L}_r$  follows  $\mathcal{L}_2$  loss, and we use a regularization loss (Lu et al. 2024)  $\mathcal{L}_S = \frac{1}{N} \sum_\xi \text{Prod}(\mathbf{S}_\xi)$  to inhibit the enlargement of scale, where  $N$  is the total number of Gaussians.

### Controllable LiDAR Generation

In this section, we introduce a controllable LiDAR-to-LiDAR generation model designed to generate geometrically authentic extrapolated scans, serving as supplementary supervision for expansive reconstruction. A primary challenge in generative novel view synthesis lies in ensuring coherence between generation and existing knowledge. Existing LiDAR generation models that use cross-modal prompt, such as semantic maps and bounding boxes (Ran, Guizilini, and Wang 2024; Pan et al. 2025), are inadequate for generating geometrically consistent scans of complex outdoor scenes, because these prompts are too sparse compared to LiDAR scans. To address this issue, we propose utilizing coarsely extrapolated LiDAR range views – rendered from

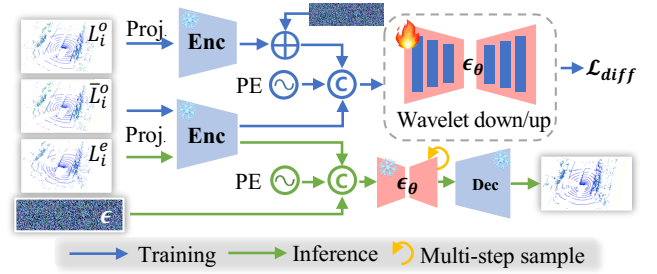


Figure 4: Workflow of the Controllable LiDAR Generation.

neural 2D Gaussian representations acquired during the initial stage of single-pass reconstruction – as conditional priors for producing LiDAR scans. This approach not only establishes a robust global geometric foundation but also unifies the modalities of the condition and output, both represented as LiDAR range views.

**Constructing Training Pair.** Due to the absence of extrapolated ground truth in single-pass driving clips, direct supervision of our controllable LiDAR generation model on extrapolated viewpoints is not feasible. Hence, we simulate low-quality GS representations to construct training pairs. Specifically, given neural 2D Gaussians reconstructed from a single-pass driving clip consisting of  $F$  frames of LiDAR scans  $\{L_i^o\}_{i=1}^F$ , we construct low-quality LiDAR rendering  $\{\bar{L}_i^o\}_{i=1}^F$  as conditional data at the reconstructed viewpoint by introducing perturbations with a variance of  $\sigma = 0.2$  to the inputs of the neural 2D Gaussian field and randomly dropout Gaussian primitives at a ratio of  $\tau = 0.1$ . In total, approximately 27k training pairs are constructed from Waymo Open dataset (Sun et al. 2020) and the Para-Lane dataset (Ni et al. 2025b), taking 20h on  $4 \times \text{RTX3090}$ .

**Training Diffusion Model.** We choose the Latent LiDAR Diffusion Model (LiDM) (Ran, Guizilini, and Wang 2024) as our base model, and introduce two additional lightweight modules to improve generation details, including adding fourier Position Encoding (PE) at input (Hu, Zhang, and Hu 2024) and integrating wavelet transforms into the up-sampling and downsampling layer (Yang et al. 2023). As shown in Fig. 4, the LiDAR’s input and condition are respectively encoded into latent space as  $z^L$  and  $c^L$  with a pre-trained VAE encoder. After diffusing the input  $z^L$  to  $z_t^L = \sqrt{\bar{\alpha}_t} z_0^L + \sqrt{1 - \bar{\alpha}_t} \epsilon$  with noise  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  and a noise schedule  $\bar{\alpha}$  at diffusion time step  $t$ , we concatenate them as the input of denoised model  $\epsilon_\theta$  and optimize it by the loss function:

$$\mathcal{L}_{diff} = \mathbb{E}_{z_0^L, \epsilon, c^L, t} \left[ \|\epsilon - \epsilon_\theta(z_t^L, t, c^L)\|_2^2 \right]. \quad (7)$$

This design enables recovery of realistic LiDAR from low-quality cues and generalizes to produce scans from extrapolated coarse rendering.

**Generated LiDAR Scan on View Extrapolation.** Given a set of 2D Gaussians  $\xi$  reconstructed from a single-pass driving clip, we render the low-quality LiDAR scans  $\{L_i^e\}_{i=1}^K$  across lanes by extrapolating the viewpoint. These

scans serve as conditional guidance for the frozen LiDAR generation model to produce geometrically consistent extrapolated scans, which provide additional supervision to refine the neural 2D Gaussians  $\xi$ .

### Depth Distortion-Aware Distillation Strategy

Even though the generative model can eliminate most of the incorrect LiDAR rendering, it still introduces hallucinations in the details (See Table 2 for ablation experiments of expansive reconstruction without DDAD). As shown in the Fig. 5 (c), we use the Chamfer Distance (CD) to quantify the difference between the generative results and the reference real scans on the Para-Lane dataset, revealing that there are still detail discrepancies between them. In fact, the Fig. 5 (b) has demonstrated that neural 2D Gaussians have preliminary ability to render global geometry in extrapolated viewpoints, with artifacts and rendering collapses occur in under-fitted regions due to insufficient extrapolated supervision. As shown in Table 2, if we fully inject generated data, the upper bound for novel views is limited by the quality of the generative data (as Fig. 5 (c)), and negative effects may also be introduced to already well-fitted regions.

Therefore, we propose a depth distortion-aware distillation strategy specifically designed to identify and correct under-fitted regions in extrapolated views, thereby minimizing the adverse impact of the LiDAR generative model for GS expansive reconstruction. Specifically, during the rendering phase, we record the median depth  $d_m$  for each ray, approximated by the depth at the Gaussian intersection where the transmittance  $T$  is closest to 0.5 (Barron et al. 2022; Huang et al. 2024). Following Eq. (5), transmittance  $T$  is defined as follows:

$$T = \prod_{j=1}^{i-1} (1 - \alpha_{\xi_j} G_{\xi_j}). \quad (8)$$

It is well-established that when transmittance approaches 0, the ray is considered to have passed through the physical surface, and rendering halts, outputting the rendered depth  $\bar{d}$ . Generally, Gaussian distributions that converge to the real depth contribute significantly to the rendering when their cumulated opacity close to 1. In these cases, as indicated by Eq. (5), the discrepancy between the median depth  $d_m$  and rendered depth  $\bar{d}$  is minimal. Conversely, a large difference suggests that the attributes of the Gaussians involved in rendering remain under-fitted. Therefore, we define areas with significant depth discrepancies as distortion regions:

$$M = \{|\bar{d}_m - \bar{d}| > \delta\}, \quad (9)$$

where  $\delta = \text{median}\{\max(s_u, s_v)\}$  is defined as the median value of the scale coefficients from the longest axis of all Gaussian trained from single traverse.

In Fig. 5 (b), this definition effectively labels ambiguous rendered point clouds as distorted regions  $M$ , allowing us to further fine-tune the Gaussian in a targeted manner. Specifically, when mixing the generated extrapolated scans as additional supervision to further expansive reconstruction(1:1 with real), the loss function is following  $\mathcal{L}_e = M \cdot \mathcal{L}$ .




Figure 5: Qualitative and quantitative comparisons of our coarse-to-fine cross-lane novel view synthesis on the Para-Lane dataset. (a) Real LiDAR scan as a reference from another-pass; (b) Rendered without diffusion priors, and the red point indicates the distorted area  $M$ ; (c) Generative result of controllable LiDAR diffusion model conditioned by coarsely extrapolated rendering; (d) Rendering after expansive reconstruction with diffusion priors.

## Experiment

### Experiment Setup

**Dataset.** We conduct experiments on both the Para-Lane (Ni et al. 2025b) and the Waymo Open (Sun et al. 2020) datasets. The Para-Lane dataset provides synchronized cross-lane LiDAR data, with a resolution of  $32 \times 1800$  at 10 Hz. The Waymo dataset provides LiDAR data along a single lane with a resolution of  $64 \times 2650$  at 10 Hz. We selected 3 clips from the Para-Lane dataset and 4 clips from the Waymo dataset to validate our method against publicly available approaches for novel view synthesis, including both viewpoint interpolation and extrapolation.

**Baseline.** We compare our method with LiDAR4D (Zheng et al. 2024), GS-LiDAR (Jiang et al. 2025), LiDAR-RT (Zhou et al. 2025) and LiDAR-GS (Chen et al. 2024b). LiDAR4D, to our best knowledge, is currently the best NeRF-based LiDAR synthesis method. Meanwhile, the remaining methods are the latest advancements in real-time synthesis based on GS. Please refer to our *supplementary material* for detailed configurations.

**Implementation details of Gaussian representation.** LiDAR-GS++ was trained for 7,000 iterations on RTX3090, with the first 5,000 iterations dedicated to reconstruction from single traverse and the last 2,000 iterations to fine reconstruction from mixed generation and real data. We randomly initialized 500,000 GS anchors from trained scan input and adopted the same growth and split strategy as Scaffold-GS (Lu et al. 2024) and the same accumulated gradient strategy as AbsGS (Ye et al. 2024). During densification, we enable it starting from step 500 with the split gradient threshold set to 0.002. We use the Adam optimizer and set the learning rate for each optimizable parameter  $\{\mathcal{F}_\Sigma, \mathcal{F}_\rho, \mathcal{F}_r, \mathcal{F}_\alpha, \mathbf{v}_\xi\}$  to  $\{1, 4, 4, 2, 5\} \times 1e^{-3}$ , respectively. For dynamic objects, we perform instance decomposition based on NSG (Ost et al. 2021) and reconstruct them separately.

**Implementation details of controllable LiDAR generation model.** The diffusion module with a frozen VAE is trained once (not per-scene) on constructed pairs for  $5e^4$  iterations using  $8 \times A100$  GPUs and batch size of 16 on each

Method	Para-Lane								Waymo								Cost	
	Lane Extrapolation @3.5m				Lane Interpolation				Lane Extrapolation @3.5m				Lane Interpolation				Train	Inference
	CD↓	F-scr↑	PSNR↑	SSIM↑	CD↓	F-scr↑	PSNR↑	SSIM↑	FRID↓	FPVD↓	JSD↓	MMD↓	CD↓	F-scr↑	PSNR↑	SSIM↑	(min)↓	(fps)↑
LiDAR4D	1.518	0.785	29.464	0.833	0.112	0.921	31.631	0.866	48.503	52.651	0.587	9.726	0.061	0.964	32.434	0.970	426	1.7
LiDAR-RT	0.482	0.806	30.430	0.824	0.159	0.902	30.979	0.837	41.330	57.551	0.576	11.738	0.063	0.962	29.410	0.917	213	<b>20.7</b>
GS-LiDAR	0.305	0.843	29.279	0.802	0.086	0.927	29.654	0.805	31.967	78.84	0.534	8.263	0.057	0.971	31.732	0.952	129	10.8
LiDAR-GS	0.270	0.865	30.742	0.828	0.090	0.923	32.103	0.869	39.095	34.018	0.606	13.301	0.059	0.968	32.313	0.973	<b>18</b>	15.8
<b>LiDAR-GS++</b>	<b>0.102</b>	<b>0.923</b>	<b>31.843</b>	<b>0.871</b>	<b>0.079</b>	<b>0.936</b>	<b>32.232</b>	<b>0.873</b>	<b>11.669</b>	<b>15.134</b>	<b>0.379</b>	<b>2.618</b>	<b>0.052</b>	<b>0.972</b>	<b>32.463</b>	<b>0.975</b>	26	16.2

Table 1: Average quantitative results on the Para-Lane dataset and Waymo Open dataset. The MMD metric is at the  $1e-5$  scale.

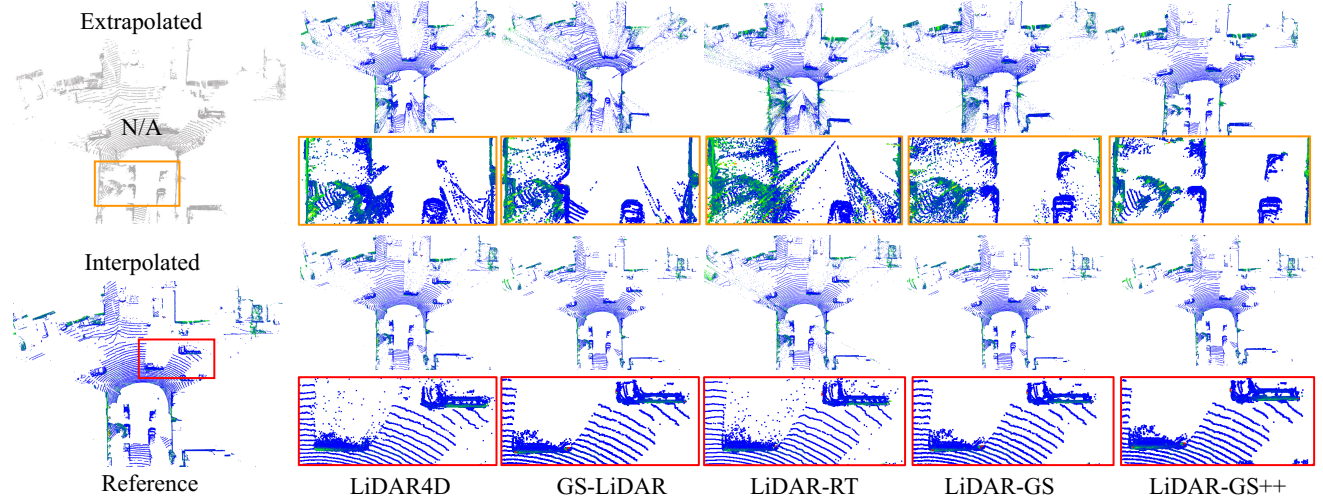


Figure 6: Qualitative comparison on Waymo sequences. Points are colorized through their intensity.

GPU. We use the Adam optimizer and set the learning rate to  $5e^{-5}$ . During inference phase, we use the DDIM (Song, Meng, and Ermon 2021) sampler with 50 sampling steps.

## Main Results

**Quantitative comparison.** On the Para-Lane dataset, we sample every 10th viewpoint from the input of a single trajectory as interpolation evaluation, and use the remaining scans as training inputs. Since this dataset provides synchronized ground-truth LiDAR scans from adjacent lanes, we use viewpoints from neighboring lanes for evaluating extrapolation. Following previous work (Chen et al. 2024b), we use CD and F-Score as evaluation metrics for point clouds, and employ PSNR and SSIM as evaluation metrics for LiDAR intensity and ray-drop.

On the Waymo dataset, the evaluation for interpolation is conducted in the same manner as described above. The difference is that, due to the lack of extrapolated ground truth, we artificially shift the input viewpoints laterally (@3.5m) as extrapolated viewpoints, and adopt the generative evaluation metrics as LiDM (Ran, Guizilini, and Wang 2024), including Fréchet Range Image Distance (FRID), Fréchet Point-based Volume Distance (FPVD), Minimum Matching Distance (MMD), Jensen-Shannon Divergence (JSD). Refer to *supplementary materials* for details.

As shown in Table 1, LiDAR-GS++ achieves excellent performance in interpolation evaluation while significantly outperforming previous methods in extrapolation evaluation. We owe this superior performance to extrapolation supervision from the diffusion prior and the expansive reconstruction of distorted areas, which enable a better representation of the neural 2DGS field. Meanwhile, our method achieves suboptimal computing efficiency. In terms of training speed, LiDAR-GS++ is only slower than LiDAR-GS due to additional computational overhead incurred during expansive reconstruction. Regarding inference speed, LiDAR-GS++ is second only to LiDAR-RT, which benefits from OptiX’s (Parker et al. 2010) hardware acceleration, enabling faster ray tracing on specific hardware configurations.

**Qualitative comparison.** Fig. 6 shows the visual comparison of LiDAR rendering. In detail, others perform poorly under extrapolated viewpoints, exhibiting numerous floater and incorrect geometry, whereas our method maintains stable and reliable performance for extrapolated re-simulation.

## Ablation Results

**Ablation on components of LiDAR-GS++.** We firstly replaced our used neural 2DGS fields (NGF) with the vanilla 2DGS (Huang et al. 2024). The results presented in Table 2

Method	Lane Extrapolation @3.5m				Lane Interpolation			
	CD↓	F-scr↑	PSNR↑	SSIM↑	CD↓	F-scr↑	PSNR↑	SSIM↑
GS-LiDAR (w/ Diff&DDAD)	0.116	0.921	29.530	0.816	0.087	0.925	29.654	0.804
Ours w/o NGF	0.417	0.825	29.878	0.809	0.095	0.922	30.039	0.816
Ours w/o Diff	0.264	0.869	30.777	0.840	0.079	<b>0.937</b>	<b>32.235</b>	0.873
Ours w/o DDAD	0.163	0.905	30.701	0.839	0.085	0.927	32.013	0.854
<b>Ours</b>	<b>0.102</b>	<b>0.923</b>	<b>31.843</b>	<b>0.871</b>	<b>0.079</b>	0.936	32.232	<b>0.873</b>

Table 2: Ablation study on Para-Lane dataset.

Method	FRID↓	FPVD↓	JSD↓	MMD↓( $1e^{-5}$ )
CLG w/o PE	33.62	15.17	0.51	13.76
CLG w/o Wavelet	29.31	14.88	0.50	10.37
CLG + BBox + Map Condition	46.74	109.51	0.66	17.72
CLG + Render Condition (Ours)	<b>28.39</b>	<b>14.46</b>	<b>0.48</b>	<b>6.20</b>

Table 3: Ablation study on Controllable LiDAR Generation.

indicate that the neural 2DGS field achieves superior reconstruction performance compared to previous Gaussian-based methods. Moreover, as shown in Table 1, our approach outperforms LiDAR-GS (Chen et al. 2024b) based 3DGS. These results underscore the critical role of the neural 2DGS field, as a robust GS representation serves as the essential foundation for achieving high-fidelity LiDAR scene reconstruction. This improvement is attributed to consideration of the impact of viewpoint and distance on LiDAR signals and the use of 2DGS primitives that are more advantageous for geometric reconstruction.

Then, to assess the necessity of incorporating diffusion priors, we conducted an experiment where extrapolated viewpoints were rendered without employing diffusion priors (w/o Diff) as additional supervision. Results in Table 2 demonstrate that the introduction of extrapolated scans generated by controllable LiDAR generation models, as supplementary supervision for expansive reconstruction, can significantly enhance the quality of extrapolated view.

Finally, we compared the performance between utilizing fully generated data for expansive reconstruction and applying Depth Distortion-Aware Distillation (DDAD) strategy. As shown in Table 2, the integration of the DDAD strategy yields superior results. Because the use of generated data without any careful adjustment not only limits the upper bound of extrapolated reconstruction but also introduces adverse effects on the real training viewpoints. As illustrated in Fig. 5, there are still detailed discrepancies between the generated data and the real reference. Our proposed strategy selectively optimizes the under-fitted regions during expansive reconstruction, preserving the convergent geometric details from the initial single traversal reconstruction.

**Analysis of the effectiveness of diffusion priors.** We integrated the diffusion prior into the previous sota, GS-LiDAR (Jiang et al. 2025), for expansive reconstruction,

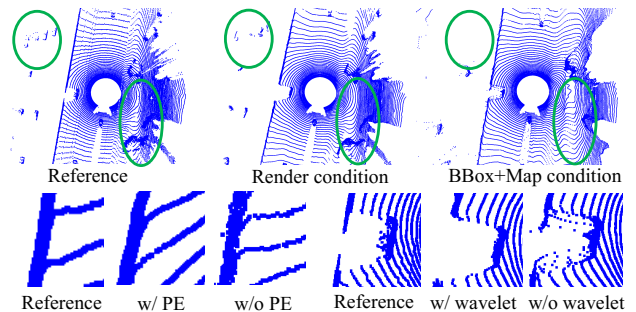


Figure 7: Qualitative visualization of CLG ablation.

which achieve substantial improvements when evaluated from extrapolation perspective, as illustrated in Table 3. These results demonstrate the effectiveness and applicability of our proposed diffusion prior for LiDAR reconstruction.

### Analysis on Controllable LiDAR Generation (CLG).

The results of the two specific ablation experiments conducted for CLG indicate that the integration of the additional modules provides positive contributions. As shown in Fig. 7, incorporating wavelet transformation aids in preserving high-frequency details during downsampling and upsampling processes, and the introduction of PE facilitates the network’s learning of continuous geometric encoding. Additionally, following (Ran, Guizilini, and Wang 2024; Pan et al. 2025), we trained a cross-modal LiDAR generation model conditioned on semantic maps and bounding boxes using the same driving clip, and compared it to our single-modal LiDAR-to-LiDAR generation model. Our model demonstrated greater potential for LiDAR synthesis tasks. Moreover, the figures reveal that although the diffusion model conditioned on semantic maps and bounding boxes shows commendable generative performance for specified prompts, it often produces geometric inconsistencies in areas lacking dense prompts, such as trees.

## Conclusion

This paper proposes LiDAR-GS++, a novel LiDAR Gaussian Splatting reconstruction framework enhanced by *diffusion priors*, designed to enable real-time and high-fidelity resimulation on public urban roads. LiDAR-GS++ leverages a controllable LiDAR generation model guided by coarsely extrapolated rendering to provide consistent additional supervision for expansive reconstruction. And we further propose a depth distortion-aware distillation strategy to efficiently refine the neural 2DGS representation in under-fitted regions. Through extensive evaluations on publicly available datasets, LiDAR-GS++ demonstrates state-of-the-art performance on both interpolated and extrapolated viewpoints.

**Limitations.** First, the reconstruction process doesn’t account for non-rigid motion of dynamic objects, e.g. pedestrians. Second, LiDAR generative models lack consideration for temporal consistency. Future work could explore the use of more advanced generation models to improve fidelity and consistency in synthesized LiDAR data.

## Acknowledgments

We thank the reviewers for the valuable discussions. This research was supported by the Zhejiang Provincial Natural Science Foundation of China under Grant No. LD24F030001.

## 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5470–5479.
- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendelevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; et al. 2023. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*.
- Cao, W.; Hallgarten, M.; Li, T.; Dauner, D.; Gu, X.; Wang, C.; Miron, Y.; Aiello, M.; Li, H.; Gilitschenski, I.; et al. 2025. Pseudo-simulation for autonomous driving. *arXiv preprint arXiv:2506.04218*.
- Chen, L.; Wu, P.; Chitta, K.; Jaeger, B.; Geiger, A.; and Li, H. 2024a. End-to-End Autonomous Driving: Challenges and Frontiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12): 10164–10183.
- Chen, Q.; Yang, S.; Du, S.; Tang, T.; Chen, P.; and Huo, Y. 2024b. LiDAR-GS: Real-time LiDAR Re-Simulation using Gaussian Splatting. *arXiv preprint arXiv:2410.05111*.
- Fan, L.; Zhang, H.; Wang, Q.; Li, H.; and Zhang, Z. 2025. Freesim: Toward free-viewpoint camera simulation in driving scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12004–12014.
- Fu, H.; Zhang, D.; Zhao, Z.; Cui, J.; Liang, D.; Zhang, C.; Zhang, D.; Xie, H.; Wang, B.; and Bai, X. 2025. ORION: A Holistic End-to-End Autonomous Driving Framework by Vision-Language Instructed Action Generation. In *IEEE/CVF International Conference on Computer Vision*, 1–18.
- Hahner, M.; Sakaridis, C.; Bijelic, M.; Heide, F.; Yu, F.; Dai, D.; and Van Gool, L. 2022. Lidar snowfall simulation for robust 3d object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16364–16374.
- Hahner, M.; Sakaridis, C.; Dai, D.; and Van Gool, L. 2021. Fog simulation on real LiDAR point clouds for 3D object detection in adverse weather. In *IEEE/CVF International Conference on Computer Vision*, 15283–15292.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- Hu, Q.; Zhang, Z.; and Hu, W. 2024. Rangeldm: Fast realistic lidar point cloud generation. In *European Conference on Computer Vision*, 115–135.
- Hu, X.; Li, S.; Huang, T.; Tang, B.; Huai, R.; and Chen, L. 2023a. How simulation helps autonomous driving: A survey of sim2real, digital twins, and parallel intelligence. *IEEE Transactions on Intelligent Vehicles*, 9(1): 593–612.
- Hu, Y.; Yang, J.; Chen, L.; Li, K.; Sima, C.; Zhu, X.; Chai, S.; Du, S.; Lin, T.; Wang, W.; Lu, L.; Jia, X.; Liu, Q.; Dai, J.; Qiao, Y.; and Li, H. 2023b. Planning-Oriented Autonomous Driving. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17853–17862.
- Huang, B.; Yu, Z.; Chen, A.; Geiger, A.; and Gao, S. 2024. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH*, 1–11.
- Huang, S.; Gojcic, Z.; Wang, Z.; Williams, F.; Kasten, Y.; Fidler, S.; Schindler, K.; and Litany, O. 2023. Neural LiDAR Fields for Novel View Synthesis. In *IEEE/CVF International Conference on Computer Vision*, 18190–18200.
- Jiang, J.; Gu, C.; Chen, Y.; and Zhang, L. 2025. GS-LiDAR: Generating Realistic LiDAR Point Clouds with Panoramic Gaussian Splatting. In *International Conference on Learning Representations*, 1–15.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4): 139:1–139:14.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*, 1–14.
- Kong, L.; Liu, Y.; Chen, R.; Ma, Y.; Zhu, X.; Li, Y.; Hou, Y.; Qiao, Y.; and Liu, Z. 2023. Rethinking range view representation for lidar segmentation. In *IEEE/CVF International Conference on Computer Vision*, 228–240.
- Li, B.; Guo, J.; Liu, H.; Zou, Y.; Ding, Y.; Chen, X.; Zhu, H.; Tan, F.; Zhang, C.; Wang, T.; et al. 2025. Uniscene: Unified occupancy-centric driving scene generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 11971–11981.
- Liao, L.; Yan, W.; Yang, M.; and Zhang, S. 2025. Learning-based 3D Reconstruction in Autonomous Driving: A Comprehensive Survey. *arXiv preprint arXiv:2503.14537*.
- Lu, T.; Yu, M.; Xu, L.; Xiangli, Y.; Wang, L.; Lin, D.; and Dai, B. 2024. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20654–20664.
- Milioto, A.; Vizzo, I.; Behley, J.; and Stachniss, C. 2019. Rangenet++: Fast and accurate lidar semantic segmentation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4213–4220. IEEE.
- Moenne-Loccoz, N.; Mirzaei, A.; Perel, O.; de Lutio, R.; Martinez Esturo, J.; State, G.; Fidler, S.; Sharp, N.; and Gojcic, Z. 2024. 3D Gaussian Ray Tracing: Fast Tracing of Particle Scenes. *ACM Transactions on Graphics*, 43(6): 1–19.
- Nakashima, K.; and Kurazume, R. 2024. Lidar data synthesis with denoising diffusion probabilistic models. In *IEEE International Conference on Robotics and Automation*, 14724–14731.
- Ni, C.; Zhao, G.; Wang, X.; Zhu, Z.; Qin, W.; Huang, G.; Liu, C.; Chen, Y.; Wang, Y.; Zhang, X.; et al. 2025a. Recondreamer: Crafting world models for driving scene reconstruction via online restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1559–1569.
- Ni, Z.; Du, S.; Hou, Z.; Wu, C.; and Yang, S. 2025b. Para-Lane: Multi-Lane Dataset Registering Parallel Scans for Benchmarking Novel View Synthesis. In *International Conference on 3D Vision*, 1–10.
- Ost, J.; Mannan, F.; Thuerey, N.; Knodt, J.; and Heide, F. 2021. Neural scene graphs for dynamic scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2856–2865.
- Pan, T.-Y.; Jeon, S.; Fan, M.; Yoo, J.; Feng, Z.; Campbell, M.; Weinberger, K. Q.; Hariharan, B.; and Chao, W.-L. 2025. Transfer Your Perspective: Controllable 3D Generation from Any Viewpoint in a Driving Scene. *arXiv preprint arXiv:2502.06682*.
- Parker, S. G.; Bigler, J.; Dietrich, A.; Friedrich, H.; Hoberock, J.; Luebke, D.; McAllister, D.; McGuire, M.; Morley, K.; Robison, A.; et al. 2010. Optix: a general purpose ray tracing engine. *ACM Transactions on Graphics*, 29(4): 1–13.

- Peng, B.; Wang, J.; Zhang, Y.; Li, W.; Yang, M.-C.; and Jia, J. 2024. Controlnext: Powerful and efficient control for image and video generation. *arXiv preprint arXiv:2408.06070*.
- Ran, H.; Guizilini, V.; and Wang, Y. 2024. Towards realistic scene generation with lidar diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14738–14748.
- Ren, X.; Huang, J.; Zeng, X.; Museth, K.; Fidler, S.; and Williams, F. 2024. Xcube: Large-scale 3d generative modeling using sparse voxel hierarchies. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4209–4219.
- Ren, X.; Lu, Y.; Cao, T.; Gao, R.; Huang, S.; Sabour, A.; Shen, T.; Pfaff, T.; Wu, J. Z.; Chen, R.; Kim, S. W.; Gao, J.; Leal-Taixe, L.; Chen, M.; Fidler, S.; and Ling, H. 2025. Cosmos-Drive-Dreams: Scalable Synthetic Driving Data Generation with World Foundation Models. *arXiv preprint arXiv:2506.09042*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising diffusion implicit models. In *International Conference on Learning Representations*.
- Sun, P.; Kretschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; et al. 2020. Scalability in perception for autonomous driving: Waymo open dataset. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2446–2454.
- Tao, T.; Gao, L.; Wang, G.; Lao, Y.; Chen, P.; Zhao, H.; Hao, D.; Liang, X.; Salzmann, M.; and Yu, K. 2024. LiDAR-NeRF: Novel lidar view synthesis via neural radiance fields. In *ACM International Conference on Multimedia*, 390–398.
- van den Oord, A.; Vinyals, O.; and Kavukcuoglu, K. 2017. Neural Discrete Representation Learning. In *International Conference on Neural Information Processing Systems*, 6309–6318.
- Wang, Q.; Fan, L.; Wang, Y.; Chen, Y.; and Zhang, Z. 2025. FreeVS: Generative view synthesis on free driving trajectory. In *International Conference on Learning Representations*, 1–23.
- Wu, H.; Zuo, X.; Leutenegger, S.; Litany, O.; Schindler, K.; and Huang, S. 2024a. Dynamic LiDAR Re-simulation using Compositional Neural Fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19988–19998.
- Wu, Y.; Zhang, K.; Qian, J.; Xie, J.; and Yang, J. 2024b. Text2lidar: Text-guided lidar point cloud generation via equirectangular transformer. In *European Conference on Computer Vision*, 291–310.
- Wu, Z.; Liu, T.; Luo, L.; Zhong, Z.; Chen, J.; Xiao, H.; Hou, C.; Lou, H.; Chen, Y.; Yang, R.; et al. 2023. Mars: An instance-aware, modular and realistic simulator for autonomous driving. In *CAAI International Conference on Artificial Intelligence*, 3–15. Springer.
- Yan, T.; Yin, J.; Lang, X.; Yang, R.; Xu, C.-Z.; and Shen, J. 2025. OLiDM: Object-aware LiDAR Diffusion Models for Autonomous Driving. In *AAAI Conference on Artificial Intelligence*, 9121–9129.
- Yan, Y.; Lin, H.; Zhou, C.; Wang, W.; Sun, H.; Zhan, K.; Lang, X.; Zhou, X.; and Peng, S. 2024. Street gaussians: Modeling dynamic urban scenes with gaussian splatting. In *European Conference on Computer Vision*, 156–173.
- Yang, X.; Zhou, D.; Feng, J.; and Wang, X. 2023. Diffusion probabilistic model made slim. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22552–22562.
- Ye, Z.; Li, W.; Liu, S.; Qiao, P.; and Dou, Y. 2024. AbsGS: Recovering fine details in 3D Gaussian Splatting. In *ACM International Conference on Multimedia*, 1053–1061.
- Zeng, X.; Du, S.; Chen, Q.; Liu, L.; Shu, H.; Gao, J.; Liu, J.; Xu, J.; Xu, J.; Chen, M.; et al. 2025. Industrial-Grade Sensor Simulation via Gaussian Splatting: A Modular Framework for Scalable Editing and Full-Stack Validation. *arXiv preprint arXiv:2503.11731*.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *IEEE/CVF International Conference on Computer Vision*, 3836–3847.
- Zhang, L.; Xiong, Y.; Yang, Z.; Casas, S.; Hu, R.; and Urtasun, R. 2024. Copilot4d: Learning unsupervised world models for autonomous driving via discrete diffusion. In *International Conference on Learning Representations*, 1–31.
- Zheng, Z.; Lu, F.; Xue, W.; Chen, G.; and Jiang, C. 2024. LiDAR4D: Dynamic Neural Fields for Novel Space-time View LiDAR Synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5145–5154.
- Zhou, C.; Fu, L.; Peng, S.; Yan, Y.; Zhang, Z.; Chen, Y.; Xia, J.; and Zhou, X. 2025. LiDAR-RT: Gaussian-based ray tracing for dynamic lidar re-simulation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1538–1548.
- Zhou, X.; Lin, Z.; Shan, X.; Wang, Y.; Sun, D.; and Yang, M.-H. 2024. Drivinggaussian: Composite gaussian splatting for surrounding dynamic autonomous driving scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21634–21643.
- Zyrianov, V.; Che, H.; Liu, Z.; and Wang, S. 2024. Lidardm: Generative lidar simulation in a generated world. *arXiv preprint arXiv:2404.02903*.
- Zyrianov, V.; Zhu, X.; and Wang, S. 2022. Learning to generate realistic lidar point clouds. In *European Conference on Computer Vision*, 17–35.