

LidarPainter: One-Step Away from Any Lidar View to Novel Guidance

Yuzhou Ji¹, Ke Ma¹, Hong Cai^{2*}, Anchun Zhang², Lizhuang Ma¹, Xin Tan^{3†}

¹Shanghai Jiao Tong University, Shanghai, China

²51WORLD, Shanghai, China

³East China Normal University, Shanghai, China

lzma@sjtu.edu.cn, caihong@51sim.com, xtan@cs.ecnu.edu.cn

Abstract

Dynamic driving scene reconstruction is of great importance in fields like digital twin system and autonomous driving simulation. However, unacceptable degradation occurs when the view deviates from the input trajectory, leading to corrupted background and vehicle models. To improve reconstruction quality on novel trajectory, existing methods are subject to various limitations including inconsistency, deformation, and time consumption. This paper proposes LidarPainter, a one-step diffusion model that recovers consistent driving views from sparse LiDAR condition and artifact-corrupted renderings in real-time, enabling high-fidelity lane shifts in driving scene reconstruction. Extensive experiments show that LidarPainter outperforms state-of-the-art methods in **speed**, **quality** and **resource efficiency**, specifically $7 \times$ faster than StreetCrafter with only one fifth of GPU memory required. LidarPainter also supports stylized generation using text prompts such as “foggy” and “night”, allowing for a diverse expansion of the existing asset library.

Proj Page — <https://george-attano.github.io/LidarPainter>

Introduction

In the fields of digital twin systems and autonomous driving simulation, the reconstruction of dynamic driving scenes is of great significance. While it is impossible to capture every single driving scenario in reality, a reconstructed scene can not only be edited to create unlimited test cases for evaluating driving algorithms at low costs, but can also unify the data collected on different sensors to build a data-consistent asset library. However, real-world driving data cannot provide the same surround view input as in static indoor scenes, causing significant degradation as the viewpoint deviates from the original trajectory. Therefore, obtaining usable reconstructed assets with comprehensive supervision for driving simulation remains a vital challenge.

To improve reconstruction quality, previous methods (Yang et al. 2024a; Ni et al. 2025; Zhao et al. 2025b; Yu et al. 2025; Zhao et al. 2025a) prove the necessity of leveraging generative priors (Gao et al. 2024; Wang et al. 2024; Zhao



Figure 1: Comparison of diffusion guidance sampled on 7000 iterations. Our method shows much better fidelity and consistency with clear characters on truck, while StreetCrafter generates corrupted vehicles and text.

et al. 2025c) as guidance. Similarly, the recent state-of-the-art method, StreetCrafter (Yan et al. 2025), trained a video diffusion model based on Vista (Gao et al. 2024), successfully generating novel view sequences with precise camera control under LiDAR conditions. However, StreetCrafter degrades largely on sparse LiDAR renderings, corrupting the geometric structure of objects and failing in most details, such as text and patterns (as shown in Figure 1), which limits the realism of reconstructed assets. StreetCrafter takes up a good deal of GPU memory, specifically 55 GB for a clip of only 15 frames. This also means a long driving sequence needs to be processed in fragments, resulting in sudden changes of scene architecture on the junction of video clips and causes artifacts. Moreover, a three-stage video diffusion sample in StreetCrafter of only 100 frames can take up to 45 minutes on a single A100 GPU, even more than the complete training time for scene reconstruction.

By analyzing existing methods, we believe that there is a common misunderstanding in improving driving scene reconstruction quality. While the consistency between image guidance and reality serves as a crucial factor in 3D scene reconstruction, many methods (Ni et al. 2025; Zhao et al. 2025b,a; Yan et al. 2025) treat video generation model as a reliable provider of novel-view guidance. However, although video generation indeed provides general consistency in the exact input sequence (e.g. the car still goes straight in the same lane), the results are not only variant concerning high-frequency information such as details of building and vehicle exteriors, but also show significant changes between stitched video clips. The detail variant nature of these methods degrades the quality of generated novel guidance and consequently limits scene reconstruction

*Project Lead

†Corresponding Author

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

results. Nevertheless, digital twin systems and autonomous driving simulation require a large asset library that needs to be obtained cost-effectively and efficiently, where the high GPU usage and slow generation time of video diffusion still largely prohibit StreetCrafter and other relevant works from facilitating real-world applications.

To get out of such a predicament, we believe the ultimate goal is to directly establish high consistency between generated images and real capture, instead of inside image sequences. Inspired by CycleGAN-turbo and Pix2Pix-turbo (Parmar et al. 2024) which achieve highly controlled image translation with consistent details, in this paper we present LidarPainter, a one-step diffusion model that effectively generates high-fidelity novel guidance from artifact rendering and LiDAR rendering for driving scene reconstruction. By transferring the original video generation task to controlled image generation, LidarPainter achieves better quality and significantly lower resource consumption, enabling real-world implementations (e.g. can be running on a single RTX 2080 Ti GPU).

Given a set of driving scene capture, we first train the 3DGS (Kerbl et al. 2023) representation for scene actors and background on early iterations with the original trajectory. Once stabilized, we render images on new trajectories, which contain many artifacts and corrupted object structures because of the lack of supervision. Meanwhile, we render LiDAR point cloud images on these novel views, which provide structural guidance. The corrupted images and paired LiDAR images are then sent through LidarPainter’s one-step diffusion model. LidarPainter introduces a Latent Attention Fusion Module that can both take in LiDAR’s structural information and preserve high frequency details, generating photorealistic images with removed artifacts and restored structure, which are made extra novel view supervision. By continuing to train the scene gaussians with both original capture and novel view guidance, we obtain a driving scene reconstruction with improved quality and view range, which can be further extended with multiple rounds of novel view expansion.

Experiments show that LidarPainter can generate high-fidelity novel view guidance with structure consistency and detail accuracy (see Figure 1), achieving more reliable trajectory interpolation and extrapolation results compared with state-of-the-art methods for driving scene reconstruction. In particular, our method is also superior in speed and resource efficiency, where LidarPainter is $7 \times$ faster than the latest StreetCrafter with only one fifth of GPU memory required. Moreover, LidarPainter further supports stylized generation using text prompts such as “foggy” and “night”, allowing for a diverse expansion of the existing asset library. Extensive ablation study shows the effectiveness of the LidarPainter network, which is promising for being adapted to other downstream tasks.

In summary, the principal contributions of this work include:

- We present LidarPainter, which to our best knowledge, is the first one-step diffusion model that can achieve better generation quality and consistency compared with video diffusion models trained on large driving datasets, en-

abling higher driving scene reconstruction quality.

- As far as we know, LidarPainter is one of the **fastest** and most **resource efficient** method for driving-scene novel-view generation, facilitating real-world applications.
- We propose the Latent Attention Fusion Module in LidarPainter, which effectively achieves region-aware fusion generation with multiple inputs, empowering controlled generation for many other downstream tasks.

Related Work

NeRF and 3D Gaussian Splatting (3DGS)

Neural Radiance Fields (NeRF) and 3D Gaussian Splatting (3DGS) have revolutionized 3D scene modeling and rendering. After first introduction, NeRF (Mildenhall et al. 2021) has advanced largely with further innovations (Kumar et al. 2025; Zheng and Xu 2025; Wu et al. 2025b). 3D Gaussian Splatting, detailed by (Kerbl et al. 2023) and expanded upon in work (Sun et al. 2025; Kwon, Cho, and Kim 2025; Lu et al. 2025) on dynamic 3D Gaussians, optimizes the rendering of Gaussian kernels for real-time applications. Other contributions (Chen et al. 2024; Tian et al. 2025b,a; Xie et al. 2026; Hu et al. 2026) also underscore the ongoing enhancements and versatility of 3DGS in handling increasingly complex rendering tasks. Noting the advantages of 3D Gaussian Splatting in rendering speed and quality, our paper chooses it as a representation for modeling dynamic driving scenes.

Driving Scene Reconstruction

Both NeRF and 3DGS can be implemented for driving scene (Tan et al. 2025) reconstruction. While there are moving traffic participants such as pedestrians and vehicles, the scene cannot be modeled the same as in static scenes.

Existing techniques on dynamic scene reconstruction are mainly divided into two categories, one is to leverage time as an additional parameter to capture temporal variations in dynamic scenes (Attal et al. 2023; Fridovich-Keil et al. 2023; Song et al. 2023; Huang et al. 2024), but these methods do not show good support for scene manipulation. The other way is to separately model scene background and moving traffic participants (Wu et al. 2023; Yang et al. 2023a,b; Tonderski et al. 2024; Yan et al. 2024; Hess et al. 2025; Chen et al. 2025), so that the scene can be easily edited, but may result in high storage and rendering overhead in large scenes without optimization.

Specifically, StreetGaussians (Yan et al. 2024) models the background and each moving object using separate Gaussian models. EmerNeRF (Yang et al. 2023a) stratifies scenes into static and dynamic fields, each modeled with a hash grid. UniSim (Yang et al. 2023b) and NeuRAD (Tonderski et al. 2024) utilize neural feature grids to model dynamic driving scenes with CNN renderer to enhance the ability of view extrapolation. But without supervision on novel trajectories, the reconstruction quality of these methods is heavily restricted.

Reconstruction With Diffusion Prior

Advancements in image and video diffusion models (Romach et al. 2022; Blattmann et al. 2023; Yang et al. 2024b)

have achieved highly controlled generation, making it possible to leverage diffusion prior as photorealistic guidance (Gong et al. 2026; Wang et al. 2026; Li et al. 2025) for scene reconstruction. DriveDreamer4D (Zhao et al. 2025a) uses driving world model (Wang et al. 2024; Zhao et al. 2025c) to directly generate novel-view videos for training, improved reconstruction quality but still has corrupted novel renderings because of poor generation consistency. StreetCrafter (Yan et al. 2025) trained a video diffusion model conditioned on LiDAR rendering to achieve precise camera control, but degrades on sparse LiDAR condition. ReconDreamer (Ni et al. 2025; Zhao et al. 2025b) trained a video restorer to recover clear views from novel view degraded images, but still suffers from high time and memory consumption, and limited quality as other video diffusion based methods.

Meanwhile, the potential of image diffusion prior is also explored. DIFIX3D+ (Wu et al. 2025a) trained a diffusion model that takes in surround view guidance to restore novel views with extreme artifacts, but creates significant ghosting areas. FreeSim (Fan et al. 2025) treats this as an image enhancement task, but can only provide acceptable quality under progressive shifting with a step size of less than 1 meter, which is impossible for real-world applications because the size of scenes grows cubically.

Method

Overview

The general pipeline of driving scene reconstruction with LidarPainter is shown in Figure 2. Given a set of driving scene capture, we first conduct an initialization train for dynamic 3DGS (Kerbl et al. 2023) representation with the original trajectory. Once stabilized, novel view images with artifacts will be rendered along with synced LiDAR point cloud images containing structural guidance, which are then sent through LidarPainter’s one-step diffusion model. LidarPainter generates photorealistic images with removed artifacts and restored structure, which are made extra novel view supervision. By continuing to train the scene gaussians with both original capture and novel view guidance, we obtain a driving scene reconstruction with improved quality and view range, which can be further extended with multiple rounds of novel view expansion.

Initialization

First, we unify all data and prepare for LiDAR rendering on novel views.

Unifying Data Driving datasets (Caesar et al. 2020; Sun et al. 2020; Xiao et al. 2021) are captured on multiple sensors including egocentric cameras, LiDAR and radars, which comes with different sampling frequency and coordinate system. Therefore, it is necessary to unify all data for the same processing paradigm to ensure temporal and spatial consistency. Specifically, LiDAR point cloud, ego poses, camera extrinsics and actor bounding box annotations are all transformed to world coordinate system. Meanwhile, we use only LiDAR information that is synchronized with camera capture to avoid deviated actor positions.

LiDAR Rendering LiDAR images are used both in training our diffusion model and scene reconstruction. In a driving scene sequence with N captured frames, the LiDAR points fall in M actor bounding box annotations build actor point cloud $\{\mathbf{P}_i^a\}_{i=0}^{M-1}$, while the rest build background point cloud $\{\mathbf{P}_i^b\}_{i=0}^{N-1}$. Given a camera pose \mathbf{C}^t on frame time $\{t_i\}_{i=0}^{N-1}$, we aggregate the background LiDAR points $\{\mathbf{P}_i^b\}_{i=0}^{N-1}$ within a temporal window of size l to form a unified point cloud \mathbf{P}^l in the world coordinate system and perform point rasterization under the given camera pose \mathbf{C}^t as in StreetCrafter (Yan et al. 2025) for background rendering. The coordinate of actor point cloud $\{\mathbf{P}_i^a\}_{i=0}^{M-1}$ built by complete capture sequence is transformed based on actor annotations on frame time $\{t_i\}_{i=0}^{N-1}$, and rendered separately.

Novel View Guidance With One-Step Diffusion

To improve scene reconstruction quality of driving scene captured on single trajectory, we leverage diffusion prior as guidance by building an efficient one-step diffusion model. We build LidarPainter upon SD-Turbo (Sauer et al. 2024), and the network structure is shown in Figure 2. In this section we introduce the design of LidarPainter, while detailed ablations can be found in Ablation Studies.

Preliminaries Diffusion Models (DM) (Sohl-Dickstein et al. 2015) are probabilistic models designed to learn a data distribution $p(x)$ by gradually denoising a normally distributed variable, which corresponds to learning the reverse process of a fixed Markov Chain of length T . For image synthesis, the models can be interpreted as an equally weighted sequence of denoising autoencoders $\epsilon_\theta(x_t, t)$, which are trained to predict a denoised variant from noisy input x_t . The simplified objective is:

$$L_{DM} = \mathbb{E}_{x, \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_\theta(x_t, t)\|_2^2 \right], \quad (1)$$

where t is uniformly sampled from $\{1, \dots, T\}$.

In Latent Diffusion Models (LDM) (Rombach et al. 2022), the data x_t is encoded into a latent representation z_t via the perceptual compression encoder \mathcal{E} , while samples from $p(z)$ can be decoded to image space with a single pass through decoder \mathcal{D} . LDM realizes the neural backbone $\epsilon_\theta(o, t)$ as a time-conditional UNet (Ronneberger, Fischer, and Brox 2015), and the simplified objective is can be:

$$L_{LDM} = \mathbb{E}_{\mathcal{E}(x), \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_\theta(z_t, t)\|_2^2 \right]. \quad (2)$$

Adding LiDAR Condition One-step diffusion has proven effective for image-to-image translation tasks (Sauer et al. 2024; Parmar et al. 2024; Wu et al. 2025a), but the task is much more demanding in generating highly consistent novel view guidance for 3D reconstruction than 2D style transfer. DIFIX3D+ (Wu et al. 2025a) first trains a single-step diffusion model to remove novel-view artifacts for expanded guidance, but fails to restore fully corrupted targets such as lane lines. Therefore, finding LiDAR renderings can provide strong structural guidance in novel views, we add LiDAR as image diffusion condition.

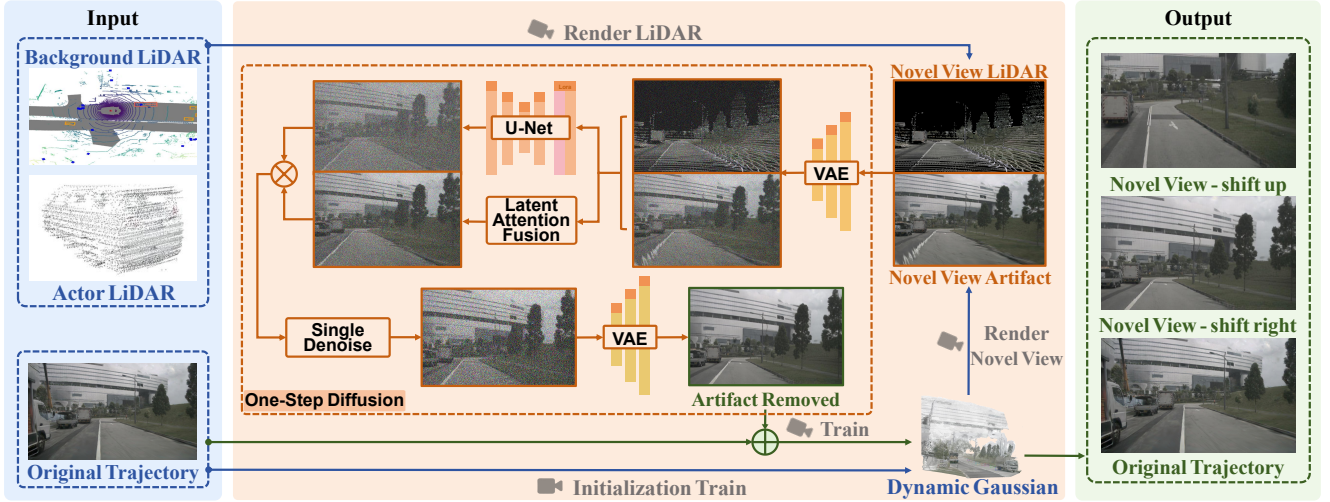


Figure 2: LidarPainter Reconstruction Pipeline.

Given an artifact-corrupted novel-view rendering I^A and corresponding LiDAR rendering I^L , we send them separately through VAE encoder \mathcal{E} to obtain latent representation z :

$$z^A = \mathcal{E}(I^A); z^L = \mathcal{E}(I^L), \quad (3)$$

which are then feed together to time-conditional UNet for predicting noise sample at timestep t :

$$Z_t^{com} = \epsilon_\theta((z^A, z^L), t). \quad (4)$$

Normally, the noise sample is denoised conditioned on original noise sample. But we find it insufficient with multiple inputs. Therefore, we propose a Latent Attention Fusion Module to achieve region-aware fusion generation.

Latent Attention Fusion While LiDAR rendering provides strong structural guidance in terms of fully corrupted targets, it is incapable of handling details such as small components and text. In order to both preserve details from artifact images and restore corruption with LiDAR rendering, we conduct pixel-wise noise fusion.

To match the fused noisy residual Z_t^{com} , we fuse latent representation z^A, z^L as:

$$z^{com} = Att(z^A, z^L) \odot z^A + (1 - Att(z^A, z^L)) \odot z^L, \quad (5)$$

where \odot denotes element-wise multiplication and Att indicates a trainable 2D convolutional module that generates pixel-wise fusion weights.

Then the predicted noise sample Z_t^{com} is denoised conditioned on fused latent z^{com} by:

$$z^D = \frac{\sqrt{\alpha_t - 1} \cdot \beta_t}{\beta_{t-1}} \cdot Z_t^{com} + \frac{\sqrt{\alpha_t} \cdot \beta_{t-1}}{\beta_t} \cdot z^{com} + \sigma_t \cdot \epsilon_t, \quad (6)$$

where α_t and β_t are the strength of original signal and added noise, ϵ_t is perturbation noise, σ_t is the standard deviation of the Gaussian noise at timestep t (controlled by the schedule).

Finally, the denoised latent z^D is decoded by \mathcal{D} to image space:

$$I^O = \mathcal{D}(z^D). \quad (7)$$

Losses We use LoRA adapters (Hu et al. 2022) in VAE and U-Net modules, and train the Latent Attention Fusion Module. We also add skip connections between the Encoder and Decoder networks as in pix2pix-turbo (Parmar et al. 2024) for fine-grained details.

For reconstruction loss, we use L2 difference and Structural Similarity Index Measure (SSIM). For perceptual loss, we use LPIPS (Zhang et al. 2018) and GAN loss (Parmar et al. 2024).

The final loss term is:

$$\mathcal{L} = 0.2\mathcal{L}_{LPIPS} + 0.6\mathcal{L}_2 + 0.4\mathcal{L}_{SSIM} + \mathcal{L}_{GAN}. \quad (8)$$

Dynamic Scene Reconstruction

In order to have strong scene editing ability, we separately model scene background and moving traffic participants for dynamic scene reconstruction.

Basic Models In LidarPainter, each Gaussian primitive \mathcal{G} is parameterized by properties $\{\mu, \Sigma, \alpha, c\}$ following the original 3DGS (Kerbl et al. 2023), where the covariance matrix Σ can be decomposed into a scaling factor $s \in \mathbb{R}_+^3$ and a rotation quaternion $r \in \mathbb{R}^4$.

While background model is already in the world coordinate system, moving objects are separately optimized as local models. Given the SE(3) pose $T_v = (R_v, t_v)$ of an vehicle v , where R_v is a 3x3 rotation matrix from SO(3) and t_v is a 3x1 translation vector in \mathbb{R}^3 , the Gaussian primitives \mathcal{G}_v can be transformed to world coordinate system by:

$$\hat{\mu}_v = R_v \mu_v + t_v, \hat{R}_v = R_v r_v^{mat}, \quad (9)$$

where $\hat{\mu}_v$ and \hat{R}_v indicate the position and rotation in world coordinate system, and r_v^{mat} denotes rotation matrix of \mathcal{G}_v .

Given a camera pose, we compute the colors by blending a set of ordered Gaussians \mathcal{N} overlapping the pixel as in 3DGS (Kerbl et al. 2023):

$$\mathbf{C} = \sum_{i \in \mathcal{N}} \mathbf{c}_i a_i \prod_{j=1}^{i-1} (1 - a_j), \quad (10)$$

where a_i is given by evaluating a 2D Gaussian with covariance Σ multiplied with a learned per-Gaussian opacity.

Reconstruction Process We first train the background and object Gaussians on the original trajectory to stabilize their 3D representations until sample iteration T_s . Then we render novel view images and use LidarPainter’s one-step diffusion to generate novel guidance. We continue training the Gaussians to final iteration T_e using both input images and novel guidance, with a sampling probability of p_{novel} for novel guidance. The process of novel view expansion can be repeated for progressively improving novel-view quality or keeping increasing view range.

Losses For reconstruction loss, we use L1 difference and Structural Similarity Index Measure (SSIM). For perceptual loss, we use LPIPS. Extra loss \mathcal{L}_g for sky mask and moving objects regularization (Yan et al. 2025) is also applied. The loss term for original trajectory is:

$$\mathcal{L}_{original} = \mathcal{L}_1 + 0.2\mathcal{L}_{SSIM} + \mathcal{L}_{LPIPS} + \mathcal{L}_g, \quad (11)$$

and novel view loss is:

$$\mathcal{L}_{novel} = \lambda_{novel}(\mathcal{L}_1 + 0.1\mathcal{L}_{SSIM} + \mathcal{L}_{LPIPS} + \mathcal{L}_g). \quad (12)$$

Experiments

Experimental Setup

In this section, we conduct extensive experiments to prove the superiority and rationality of LidarPainter design.

Datasets and Metrics For quantitative experiments, we train and evaluate the models on Waymo Open Dataset (Sun et al. 2020), PandaSet (Xiao et al. 2021) and nuScenes dataset (Caesar et al. 2020). Our one-step diffusion model is trained on 100 sequences of Waymo Open Dataset and nuScenes dataset including 5000 frames. We use 15 sequences from Waymo, 5 sequences from PandaSet (as in StreetCrafter (Yan et al. 2025)) and 10 sequences from nuScenes to test novel view synthesis ability.

To evaluate the quality of scene reconstruction, we use the peak signal-to-noise ratio (PSNR) and perceptual distance (LPIPS) (Zhang et al. 2018). To evaluate the quality of novel view generation, we use fr chet inception distance (FID) (Heusel et al. 2017).

Implementation Details LidarPainter’s diffusion model is trained for 36 hours using Adam (Kingma and Ba 2014) optimizer, with each image resized to 512×512 . During inference, the generated image size is 1024×576 . For original trajectory, we uniformly sample half of the images in each sequence as the testing frames and use the remaining for training the same as in StreetCrafter. We test LidarPainter with single round of novel sample. Sample iteration T_s and final iteration T_e for training are set to 7000 and 30,000. Sampling probability p_{novel} for novel guidance is 0.4, and λ_{novel} is 0.2. We use the same λ_{novel} and sampling probability p_{novel} for LidarPainter, StreetCrafter and DIFIX during experiments, and others all follow their original settings. We use the same artifact and LiDAR rendering input for LidarPainter and StreetCrafter, while providing DIFIX with ground-truth training views as reference. All results (including the training for diffusion model) are reported running on a single NVIDIA A100 80GB GPU.

Methods	Interpolation		Lane Shift @ FID↓		
	PSNR↑	LPIPS↓	side 2m	side 3m	up 1.5m
EmerNeRF (2023)	26.09	0.199	90.17	109.14	87.73
StreetGaussians (2024)	<u>30.82</u>	0.146	73.96	94.68	71.43
StreetCrafter (2025)	30.01	0.072	57.12	69.35	58.04
DIFIX (2025)	29.13	0.082	58.17	77.80	63.19
LidarPainter (Ours)	31.04	0.067	53.71	67.52	52.53

Table 1: Quantitative results on Waymo Open Dataset.

Methods	Interpolation		Lane Shift @ FID↓		
	PSNR↑	LPIPS↓	side 2m	side 3m	up 1.5m
NeuRAD (2024)	27.24	0.109	73.62	95.40	70.31
StreetGaussians (2024)	27.56	0.157	77.84	96.32	73.73
SplatAD (2025)	28.02	0.099	69.25	87.62	68.83
StreetCrafter (2025)	27.16	0.097	64.34	79.01	61.03
DIFIX (2025)	26.98	0.101	65.82	86.31	66.69
LidarPainter (Ours)	<u>27.83</u>	0.072	60.32	69.41	57.33

Table 2: Quantitative results on nuScenes Dataset.

Comparison With Other State-of-the-art Methods

We compare LidarPainter with other state-of-the-art methods by evaluating the rendered quality of trajectory interpolation and extrapolation in reconstructed scenes. We report PSNR and LPIPS on trajectory interpolation, and evaluate FID on trajectory extrapolation since there is no ground truth for novel views. We compare with EmerNeRF (Yang et al. 2023a), UniSim (Yang et al. 2023b) and NeuRAD (Tonderski et al. 2024) (NeRF-based methods), StreetGaussians (Yan et al. 2024), StreetCrafter (Yan et al. 2025) and SplatAD (Hess et al. 2025) (3DGS-based methods). We also compare with the DIFIX3D+ (Wu et al. 2025a), a static scene reconstruction method equipped with the first-tier artifact removal module, by replacing LidarPainter with DIFIX in our reconstruction pipeline.

Qualitative Comparison Qualitative visual results of image generation and 3D scene reconstruction are shown in Figure 3, with framed parts for comparison.

For the task of generating novel view guidance with artifact rendering, the latest diffusion methods StreetCrafter and DIFIX both destroy text regions and signs. StreetCrafter also generates corrupted vehicles and even cars that are in-existent in the input rendering, while DIFIX fails to restore deformed lane lines (even with ground truth view as reference). Only our method can simultaneously fix lane lines, remove artifacts and preserve details.

For the task of 3D scene reconstruction, StreetGaussian solely trained on original trajectory creates highly deformed

Methods	Interpolation		Lane Shift @ FID↓		
	PSNR↑	LPIPS↓	side 2m	side 3m	up 1.5m
UniSim (2023)	25.61	0.121	74.34	94.17	74.11
NeuRAD (2024)	27.19	0.106	63.45	87.42	61.51
StreetGaussians (2024)	27.40	0.117	69.84	89.31	66.73
SplatAD (2025)	27.51	0.103	67.82	87.42	65.01
StreetCrafter (2025)	27.39	0.098	66.74	80.21	63.96
LidarPainter (Ours)	<u>27.50</u>	0.086	61.16	77.35	60.75

Table 3: Quantitative results on PandaSet.

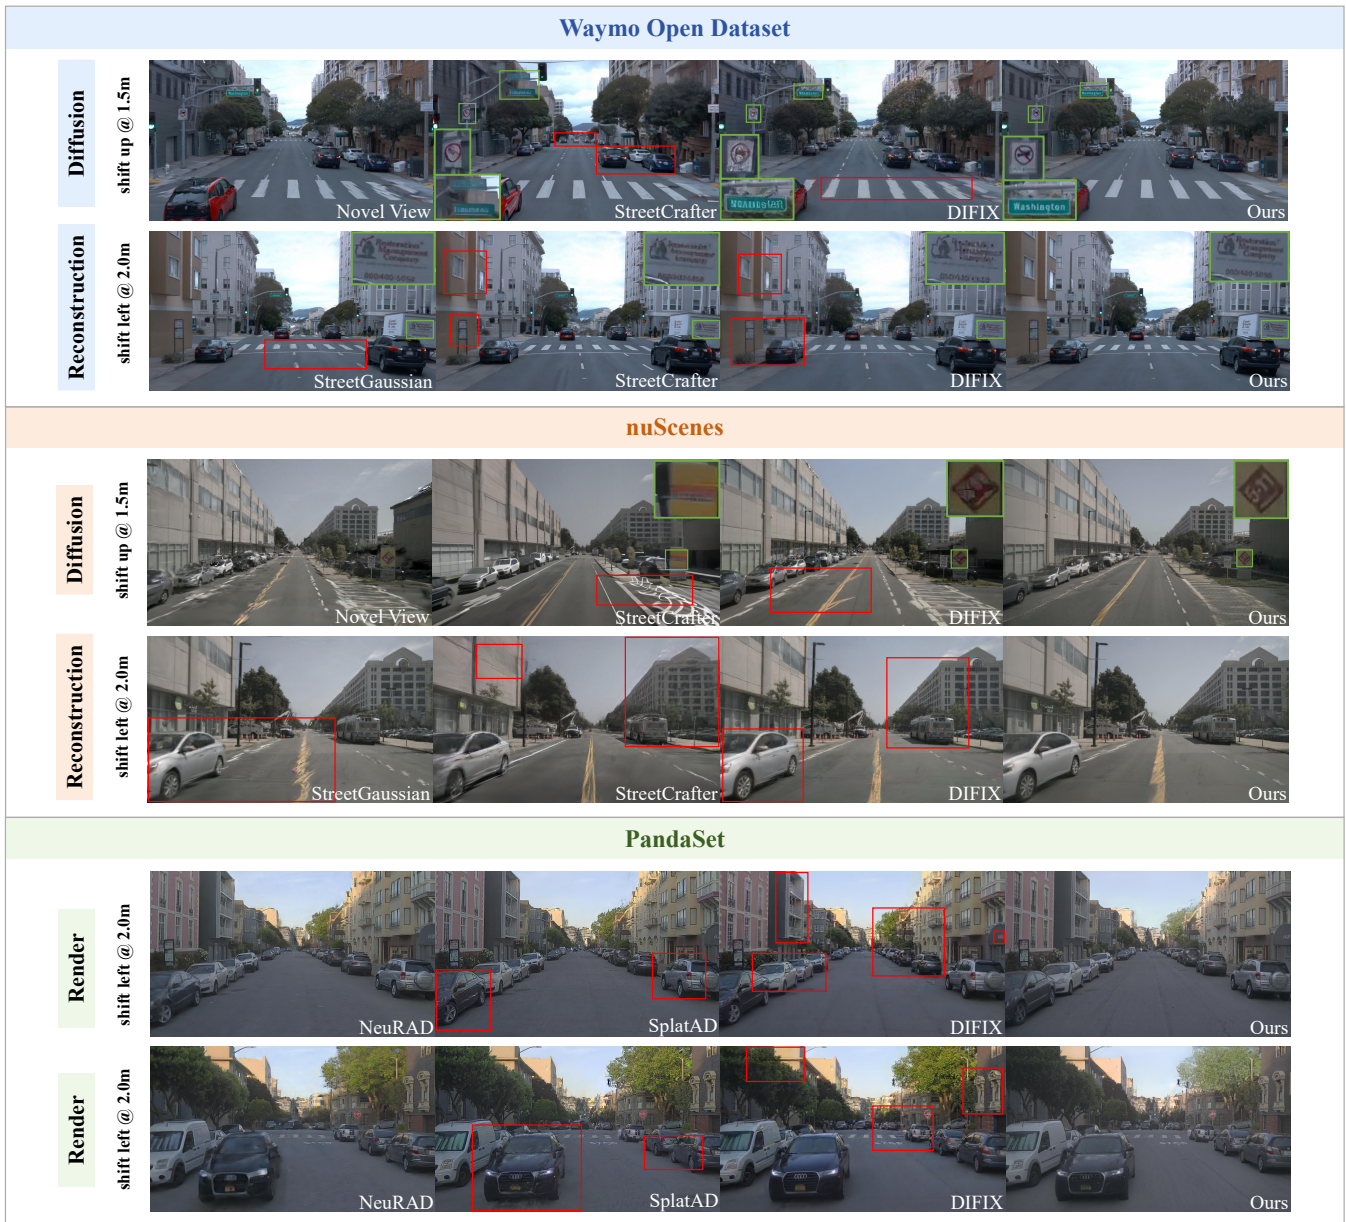


Figure 3: Qualitative comparisons of image generation and 3D scene reconstruction results on different scenes.

lane lines in novel trajectory due to the lack of supervision. With inconsistent novel guidance, StreetCrafter and DIFIX still fail in text regions, resulting in new artifacts on vehicles and buildings, while our method provides the smoothest reconstructed results on novel trajectory.

We also test LidarPainter by implementing it for post render processing at rendering stage of reconstructed scenes as in DIFIX3D+ (Wu et al. 2025a). In Figure 3, results on PandaSet shows that both DIFIX and LidarPainter can reduce artifacts comparing with single NeRF and 3DGS rendering. However, DIFIX still generates corrupted structure and transparent artifact, which contradicts the original intention of using post-processing to improve image quality.

Quantitative Comparison We report quantitative results of trajectory interpolation (original trajectory) and extrapolation (lane shift, novel trajectory) in Table 1, 2 and 3.

In trajectory interpolation, although methods such as StreetGaussians and SplatAD are solely trained on original trajectory using ground truth training views (with around 10000 more iterations), our method can still have competitive results with even better PSNR on Waymo Open Dataset. This means that correct novel guidance will improve the reconstruction quality on the original trajectory.

In trajectory extrapolation, we test the novel view quality on side (left and right) and up shifted cameras. Novel guidance is important in reducing artifacts and restore cor-

Method	StreetCrafter	DIFIX*	OURS
Inference (s/frame)	6.27	1.06	0.87
GPU Usage (GB)	55.3	10.9	10.5

Table 4: Comparison of performance and time consumption computed on rendering resolution 1024×576 (*only single shifting time in progressive update).

Input	Ablations				
Artifact			✓	✓	✓
LiDAR	✓	✓			✓
GT	✓		✓		
FID↓	109.72	96.94	79.63	73.31	66.84

Table 5: Ablation on input guidance. Artifact, LiDAR and GT denotes artifact rendering, LiDAR rendering and ground truth training view.

ruptions, thus methods trained only on original trajectory degrades significantly. Meanwhile, LidarPainter generates more consistent novel supervision than StreetCrafter and DIFIX, resulting in the lowest FID.

Speed And Resource Efficiency To obtain abundant 3D assets for autonomous driving simulation systems, the speed and resource efficiency for a reconstruction pipeline are equally important as its quality. While reconstruction module can be easily changed, we compare the speed and consumption of LidarPainter with StreetCrafter’s video diffusion and DIFIX3D+’s image diffusion. The result is shown in Table 4, where LidarPainter shows clear advantage.

In conclusion, LidarPainter is superior in quality, speed, and resource efficiency as mentioned above, showing more possibility in real-world applications.

Ablation Studies

We conduct ablation to validate the effectiveness of LidarPainter design.

Input Guidance We analyze the influence of different input guidance for LidarPainter on view generation quality. The comparison of novel view FID is reported in Table 5. Only using artifact rendering for training will result in a de-blur model that can successfully restored blurred areas but cannot restore corrupted structure, while single LiDAR input maintains structure yet provides poor image quality. Although the intention is to use the best guidance, adding nearby ground truth image as reference (similar to DIFIX3D+ (Wu et al. 2025a)) will lead to even worse results with extra artifacts, because the model cannot mix the different views properly without spatial guidance. Use paired (same view) artifact rendering and LiDAR rendering can easily use both guidance and generate better images.

Latent Attention Fusion We analyze the necessity of our latent attention fusion module. As shown in Figure 4, without Latent Attention Fusion module, the diffusion process is more of a de-blur process that removes surrounding ar-



Figure 4: Ablation on Latent Attention Fusion (LAF).




Figure 5: Prompted generation of LidarPainter.

tifacts, but fails in restoring torn lane lines. The Latent Attention Fusion module can correctly use LiDAR guidance to restore highly corrupted structure and preserve its original de-blur ability for other regions with high-frequency information such as text and vehicles.

Stylized Generation

LidarPainter also supports stylized generation based on text prompt by conditioning the latent representation z on CLIP (Radford et al. 2021) embeddings. We trained another LidarPainter model using 600 paired images of original capture, foggy and night scenes (generated by CycleGAN-turbo (Parmar et al. 2024)) with corresponding text prompts. As shown in Figure 5, with the same input LiDAR rendering, LidarPainter can generate views on different circumstances according to text prompts and maintain the original structure given by LiDAR. This allows for a diverse expansion of the existing asset library without additional real-world capture, which could be expensive and time-consuming.

Conclusion

In this paper, we present LidarPainter, a one-step diffusion model that recovers consistent driving views from sparse LiDAR condition and artifact-corrupted renderings, enabling high-fidelity lane shifts in dynamic street scene reconstruction. We present Latent Attention Fusion module implemented in one-step diffusion model to effectively use both structural information in LiDAR and details in artifact rendering, providing a promising solution for controlled image generation with multiple inputs. LidarPainter is superior in speed, quality, and resource efficiency, which also supports stylized generation using text prompts, allowing for a diverse expansion of the existing asset library.

Acknowledgments

This work is done during an internship at 51WORLD, and is supported by the National Natural Science Foundation of China No.62302167, 62472282, 72192821, U23A20343; Shanghai Sailing Program 23YF1410500; Young Elite Scientists Sponsorship Program by CAST YESS20240780; the Chenguang Program of Shanghai Education Development Foundation and Shanghai Municipal Education Commission 23CGA34; Natural Science Foundation of Chongqing CSTB2023NSCQ-MSX0137; the Fundamental Research Funds for the Central Universities (project number: YG2023QNA35).

References

- Attal, B.; Huang, J.-B.; Richardt, C.; Zollhoefer, M.; Kopf, J.; O'Toole, M.; and Kim, C. 2023. HyperReel: High-fidelity 6-DoF video with ray-conditioned sampling. In *Proc. IEEE CVPR*.
- 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*.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuscenes: A multimodal dataset for autonomous driving. In *Proc. IEEE CVPR*.
- Chen, Y.; Xu, H.; Zheng, C.; Zhuang, B.; Pollefeys, M.; Geiger, A.; Cham, T.-J.; and Cai, J. 2024. MVSplat: Efficient 3D Gaussian Splatting from Sparse Multi-View Images. In *European Conference on Computer Vision (ECCV)*.
- Chen, Z.; Yang, J.; Huang, J.; de Lutio, R.; Esturo, J. M.; Ivanovic, B.; Litany, O.; Gojcic, Z.; Fidler, S.; Pavone, M.; Song, L.; and Wang, Y. 2025. OmniRe: Omni Urban Scene Reconstruction. In *The Thirteenth International Conference on Learning Representations*.
- Fan, L.; Zhang, H.; Wang, Q.; Li, H.; and Zhang, Z. 2025. Freesim: Toward free-viewpoint camera simulation in driving scenes. In *Proc. IEEE CVPR*, 12004–12014.
- Fridovich-Keil, S.; Meanti, G.; Warburg, F. R.; Recht, B.; and Kanazawa, A. 2023. K-planes: Explicit radiance fields in space, time, and appearance. In *Proc. IEEE CVPR*.
- Gao, S.; Yang, J.; Chen, L.; Chitta, K.; Qiu, Y.; Geiger, A.; Zhang, J.; and Li, H. 2024. Vista: A generalizable driving world model with high fidelity and versatile controllability. *arXiv preprint arXiv:2405.17398*.
- Gong, J.; Zhang, C.; Liu, F.; Fan, K.; Zhou, Q.; Tan, X.; Zhang, Z.; and Xie, Y. 2026. Diffusion Implicit Policy for Unpaired Scene-aware Motion Synthesis. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Hess, G.; Lindström, C.; Fatemi, M.; Petersson, C.; and Svensson, L. 2025. Splatad: Real-time lidar and camera rendering with 3d gaussian splatting for autonomous driving. In *Proc. IEEE CVPR*, 11982–11992.
- 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, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hu, J.; Zhang, Z.; Gong, J.; Ma, L.; Tan, X.; and Xie, Y. 2026. Multi-Step Deformable Gaussian Splatting for Dynamic Scene Rendering. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Huang, N.; Wei, X.; Zheng, W.; An, P.; Lu, M.; Zhan, W.; Tomizuka, M.; Keutzer, K.; and Zhang, S. 2024. s^3 Gaussian: Self-Supervised Street Gaussians for Autonomous Driving. *arXiv preprint arXiv:2405.20323*.
- 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.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kumar, A.; et al. 2025. DynaMoDe-NeRF: Motion-aware Deblurring Neural Radiance Field for Dynamic Scenes. In *Proc. IEEE CVPR*, 21728–21738.
- Kwon, J.; Cho, H.; and Kim, J. 2025. Efficient Dynamic Scene Editing via 4D Gaussian-based Static-Dynamic Separation. *arXiv preprint arXiv:2502.02091*.
- Li, X.; Tan, X.; Chen, Z.; Zhang, Z.; Zhang, R.; Guo, R.; Jiang, G.; Chen, Y.; Qu, Y.; and Ma, L. 2025. One-for-More: Continual Diffusion Model for Anomaly Detection. *arXiv preprint arXiv:2502.19848*.
- Lu, Y.; Zhou, Y.; Liu, D.; Liang, T.; and Yin, Y. 2025. Bardgs: Blur-aware reconstruction of dynamic scenes via gaussian splatting. In *Proc. IEEE CVPR*, 16532–16542.
- 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.
- Ni, C.; Zhao, G.; Wang, X.; Zhu, Z.; Qin, W.; Huang, G.; Liu, C.; Chen, Y.; Wang, Y.; Zhang, X.; et al. 2025. Recondreamer: Crafting world models for driving scene reconstruction via online restoration. In *Proc. IEEE CVPR*.
- Parmar, G.; Park, T.; Narasimhan, S.; and Zhu, J.-Y. 2024. One-step image translation with text-to-image models. *arXiv preprint arXiv:2403.12036*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proc. IEEE CVPR*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference*,

Munich, Germany, October 5-9, 2015, proceedings, part III 18.

Sauer, A.; Lorenz, D.; Blattmann, A.; and Rombach, R. 2024. Adversarial diffusion distillation. In *European Conference on Computer Vision*.

Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*.

Song, L.; Chen, A.; Li, Z.; Chen, Z.; Chen, L.; Yuan, J.; Xu, Y.; and Geiger, A. 2023. Nerfplayer: A streamable dynamic scene representation with decomposed neural radiance fields. *IEEE Transactions on Visualization and Computer Graphics*.

Sun, P.; Kretschmar, H.; Dotiwala, 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 *Proc. IEEE CVPR*.

Sun, S.; Zhao, C.; Sun, Z.; Chen, Y. V.; and Chen, M. 2025. SplatFlow: Self-Supervised Dynamic Gaussian Splatting in Neural Motion Flow Field for Autonomous Driving. In *Proc. IEEE CVPR*, 27487–27496.

Tan, X.; Wu, W.; Zhang, Z.; Fan, C.; Peng, Y.; Zhang, Z.; Xie, Y.; and Ma, L. 2025. GEOcc: Geometrically Enhanced 3D Occupancy Network With Implicit-Explicit Depth Fusion and Contextual Self-Supervision. *IEEE Transactions on Intelligent Transportation Systems*.

Tian, Q.; Tan, X.; Gong, J.; Xie, Y.; and Ma, L. 2025a. UniForward: Unified 3D Scene and Semantic Field Reconstruction via Feed-Forward Gaussian Splatting from Only Sparse-View Images. [arXiv:2506.09378](https://arxiv.org/abs/2506.09378).

Tian, Q.; Tan, X.; Xie, Y.; and Ma, L. 2025b. Driving-Forward: Feed-forward 3D Gaussian Splatting for Driving Scene Reconstruction from Flexible Surround-view Input. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Tonderski, A.; Lindström, C.; Hess, G.; Ljungbergh, W.; Svensson, L.; and Petersson, C. 2024. Neurad: Neural rendering for autonomous driving. In *Proc. IEEE CVPR*.

Wang, S.; Zhang, J.; Tan, X.; Xie, Z.; Wang, C.; and Ma, L. 2026. MMoFusion: Multi-modal co-speech motion generation with diffusion model. *Pattern Recognition*.

Wang, X.; Zhu, Z.; Huang, G.; Chen, X.; Zhu, J.; and Lu, J. 2024. DriveDreamer: Towards Real-World-Drive World Models for Autonomous Driving. In *European Conference on Computer Vision*.

Wu, J. Z.; Zhang, Y.; Turki, H.; Ren, X.; Gao, J.; Shou, M. Z.; Fidler, S.; Gojcic, Z.; and Ling, H. 2025a. Difix3d+: Improving 3d reconstructions with single-step diffusion models. In *Proc. IEEE CVPR*.

Wu, S.; Basu, S.; Broedermann, T.; Van Gool, L.; and Sakaridis, C. 2025b. PBR-NeRF: Inverse Rendering with Physics-Based Neural Fields. In *Proc. IEEE CVPR*, 10974–10984.

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*.

Xiao, P.; Shao, Z.; Hao, S.; Zhang, Z.; Chai, X.; Jiao, J.; Li, Z.; Wu, J.; Sun, K.; Jiang, K.; et al. 2021. Pandaset: Advanced sensor suite dataset for autonomous driving. In *2021 IEEE international intelligent transportation systems conference (ITSC)*.

Xie, Z.; Gong, J.; Tan, X.; Zhang, Z.; and Xie, Y. 2026. Zero-Shot Robotic Manipulation via 3D Gaussian Splatting-Enhanced Multimodal Retrieval-Augmented Generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

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*.

Yan, Y.; Xu, Z.; Lin, H.; Jin, H.; Guo, H.; Wang, Y.; Zhan, K.; Lang, X.; Bao, H.; Zhou, X.; et al. 2025. Streetcrafter: Street view synthesis with controllable video diffusion models. In *Proc. IEEE CVPR*.

Yang, J.; Ivanovic, B.; Litany, O.; Weng, X.; Kim, S. W.; Li, B.; Che, T.; Xu, D.; Fidler, S.; Pavone, M.; et al. 2023a. Emernerf: Emergent spatial-temporal scene decomposition via self-supervision. [arXiv preprint arXiv:2311.02077](https://arxiv.org/abs/2311.02077).

Yang, Z.; Chen, Y.; Wang, J.; Manivasagam, S.; Ma, W.-C.; Yang, A. J.; and Urtasun, R. 2023b. Unisim: A neural closed-loop sensor simulator. In *Proc. IEEE CVPR*.

Yang, Z.; Pan, Z.; Yang, Y.; Zhu, X.; and Zhang, L. 2024a. Driving Scene Synthesis on Free-form Trajectories with Generative Prior. [arXiv preprint arXiv:2412.01717](https://arxiv.org/abs/2412.01717).

Yang, Z.; Teng, J.; Zheng, W.; Ding, M.; Huang, S.; Xu, J.; Yang, Y.; Hong, W.; Zhang, X.; Feng, G.; et al. 2024b. Cogvideox: Text-to-video diffusion models with an expert transformer. [arXiv preprint arXiv:2408.06072](https://arxiv.org/abs/2408.06072).

Yu, Z.; Wang, H.; Yang, J.; Wang, H.; Cao, J.; Ji, Z.; and Sun, M. 2025. Sgd: Street view synthesis with gaussian splatting and diffusion prior. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.

Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. IEEE CVPR*, 586–595.

Zhao, G.; Ni, C.; Wang, X.; Zhu, Z.; Zhang, X.; Wang, Y.; Huang, G.; Chen, X.; Wang, B.; Zhang, Y.; et al. 2025a. Drivedreamer4d: World models are effective data machines for 4d driving scene representation. In *Proc. IEEE CVPR*.

Zhao, G.; Wang, X.; Ni, C.; Zhu, Z.; Qin, W.; Huang, G.; and Wang, X. 2025b. Recondreamer++: Harmonizing generative and reconstructive models for driving scene representation. [arXiv preprint arXiv:2503.18438](https://arxiv.org/abs/2503.18438).

Zhao, G.; Wang, X.; Zhu, Z.; Chen, X.; Huang, G.; Bao, X.; and Wang, X. 2025c. Drivedreamer-2: Llm-enhanced world models for diverse driving video generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Zheng, X.; and Xu, D. 2025. Flow-NeRF: Joint Learning of Geometry, Poses, and Dense Flow within Unified Neural Representations. In *Proc. IEEE CVPR*, 993–1002.