

# LSD-3D: Large-Scale 3D Driving Scene Generation with Geometry Grounding

Julian Ost<sup>1\*</sup>, Andrea Ramazzina<sup>2,3\*</sup>, Amogh Joshi<sup>1\*</sup>, Maximilian Bömer<sup>4</sup>,  
Mario Bijelic<sup>1,4</sup>, Felix Heide<sup>1,4</sup>

<sup>1</sup>Princeton University

<sup>2</sup>Mercedes-Benz

<sup>3</sup>Technical University of Munich

<sup>4</sup>Torc Robotics

## Abstract

Large-scale scene data is essential for training and testing in robot learning. Neural reconstruction methods have promised the capability of reconstructing large physically-grounded outdoor scenes from captured sensor data. However, these methods have baked-in static environments and only allow for limited scene control – they are functionally constrained in scene and trajectory diversity by the captures from which they are reconstructed. In contrast, generating driving data with recent image or video diffusion models offers control, however, at the cost of geometry grounding and causality. In this work, we aim to bridge this gap and present a method that directly generates large-scale 3D driving scenes with accurate geometry, allowing for causal novel view synthesis with object permanence and explicit 3D geometry estimation. The proposed method combines the generation of a proxy geometry and environment representation with score distillation from learned 2D image priors. We find that this approach allows for high controllability, enabling the prompt-guided geometry and high-fidelity texture and structure that can be conditioned on map layouts – producing realistic and geometrically consistent 3D generations of complex driving scenes.

**Project Webpage** — <https://light.princeton.edu/lsd-3d>

## 1 Introduction

Large-scale public datasets have driven significant advancements in robot learning over the last decade. For autonomous driving, large volumes of data (Geiger, Lenz, and Urtasun 2012; Sun et al. 2020; Caesar et al. 2020) have unlocked new capabilities for perception and planning. Initially limited to a few labeled scenes (Geiger, Lenz, and Urtasun 2012), existing datasets now offer thousands of multi-modal, fully annotated scenes around the world (Caesar et al. 2020; Sun et al. 2020) – unlocking broader capabilities for robot learning. However, achieving generalized robot autonomy demands even greater scales of both *diverse* and *complete* data, to capture the long tail of driving scene distributions – a challenge given the high costs of capturing and annotating data. Research in various layers of the robotics stack has demonstrated that training models on large quantities of data and

in simulated environments (Talwar et al. 2020; Gulino et al. 2023) can result in robust and more generalized autonomy.

Recent work in neural scene reconstruction promises to bridge this gap by reconstructing previously simulated driving environments directly from sensor data (Ost et al. 2021; Tonderski et al. 2023; Yang et al. 2023b). Such scenes can be replayed in real time, synthesize novel views, and allow for unseen actor variations in closed-loop testing (Ljungbergh et al. 2024; Dauner et al. 2024). However, neural reconstruction methods are fundamentally limited in that they cannot produce novel content beyond recorded scenes – they do not offer data-driven simulation with great scene diversity.

Video diffusion models have been proposed to increase data volumes and diversity by subsequently fine-tuning on autonomous driving data to mimic driving datasets and multi-view camera setups (Gao et al. 2023). While these methods can generate a large corpus of novel driving data, outside of their training data, and provide feature control, they also come with inherent limitations: High computational costs restrict real-time replay and scalability for closed-loop simulation tasks. Furthermore, they lack explicit spatial modeling, which prevents causality, object permanence, and 3D consistency. As such, video diffusion models struggle as simulators, especially for safe robot learning.

Explicit 3D scene models *guarantee causality* and 3D consistency. Directly generating explicit 3D scenes, however, poses a challenge: both *geometry* and *texture* have to be generated consistently and with high quality. Some approaches use LiDAR point clouds to produce pure geometry without detailed texture (Lee et al. 2024) – which cannot be used to train image-based autonomous driving models – while more recent approaches adopt coarse 3D geometry as a conditioning mechanism for video synthesis (Lu et al. 2024a,b). As such, these methods inherit video diffusion limitations, such as a lack of causality. Distillation methods (Poole et al. 2022) instead address the 3D data gap by transferring priors from 2D image models into 3D representations via inverse rendering techniques (Yi et al. 2024). However, existing techniques are limited to object-centric generation (Poole et al. 2022; Xiang et al. 2024) and lack realism (Li et al. 2024) — and so, they are not suitable for the complexity (Zhang et al. 2024b) or spatial scale (Shriram et al. 2024) of large-scale driving scenes. Only recent work has explored distillation approaches for scene extrapolation

\*Indicates equal contribution.

from sparse real data captures (Wu et al. 2025; Liu, Zhou, and Huang 2024; Fischer et al. 2025), achieving improved reconstruction quality from few images and hinting at the potential for complete outdoor scene generation.

We propose a novel approach that overcomes the aforementioned limitations of existing generative methods for large-scale driving scenes. Our method generates explicit 3D models of entirely novel environments – with both geometry and high-fidelity texture – by fusing the diversity of image and video generation with the efficiency of explicit 3D representations. We first generate a coarse geometry of street scenes, optionally controlled by a road map layout. This proxy is then used to guide the generation of fine structural details and high-fidelity textures via image space distillation with a high-quality image generation model. Specifically, we introduce geometry-grounded distillation sampling (GGDS), an image space sampling approach, that incorporates explicit geometry control and exact noise sampling by DDIM inversion in a single method. We find that the combination of geometry guidance and consistent noise sampling can deliver 3D consistent scene generation via distillation. The method produces diverse, realistic, and large-scale 3D scenes for autonomous driving. Generated scenes guarantee causality and can produce unlimited novel trajectories in real-time – enabling *scalability* – while maintaining 3D consistency and appearance fidelity. Furthermore, precise prompt control over weather, season, time-of-day, and location, in the form of explicit environmental lighting, enables further fine-grained customization of these virtual scenes.

We validate our method on the Waymo Open Dataset (Sun et al. 2020), generating novel scenes which not only inherit the data prior distribution, but leverage the implicit prior of 2D diffusion models to provide enhanced scene *diversity*. Our approach, generating complete and coherent large-scale 3D scenes, outperforms state-of-the-art generative methods in image synthesis of unseen scenes by **18% in FVD** and maintains prompt adherence on the level of video models.

We summarize our specific contributions as follows:

- To our knowledge, our method is the first to utilize a distillation approach to directly *generate* and optimize explicit 3D driving scenes with high-quality geometry and texture – guaranteeing causal generation.
- We introduce Geometry-Grounded Distillation Sampling (GGDS), a method combining controlled proxy mesh generation with a conditional diffusion prior to produce novel, view-consistent Gaussian splatting scenes, with real-time rendering and composability with 3D assets.
- We generate diverse large-scale scenes which can be rendered into physically-grounded videos controlled by trajectories through the scene, enabling the creation of unlimited, completely unseen environments, controlled by scene descriptions, traffic map layouts, or text prompts.

## 2 Related Work

**Image and Video Synthesis.** Recent advances in image generation have enabled the synthesis of high-resolution, photorealistic imagery. These approaches – from generative adversarial networks (GANs) (Goodfellow et al. 2014;

Brock, Donahue, and Simonyan 2018; Karras et al. 2020), likelihood-based methods (Kingma and Welling 2013; Razavi, Van den Oord, and Vinyals 2019), to more prominently, diffusion models (Dhariwal and Nichol 2021; Song, Meng, and Ermon 2020; Ho, Jain, and Abbeel 2020; Rombach et al. 2022) – have been recently extended to video generation (Blattmann et al. 2023; Ho et al. 2022b; Hong et al. 2022; Team 2024). Early methods offer control via text prompt (Singer et al. 2022), image (Zhang et al. 2024c), or camera position (Sargent et al. 2023; Yang et al. 2024b). Recent methods like GEN3C (Ren et al. 2025b; Team et al. 2025) or GeoDrive (Chen et al. 2025a) for driving scenes incorporate 3D conditioning and point projection to improve geometric consistency. Specialized models, such as Vista (Gao et al. 2024), MagicDrive (Gao et al. 2023), or DriveDreamer (Wang et al. 2023b), have been developed on top of foundational video models (Blattmann et al. 2023) for autonomous driving applications – employing bounding box, HD-map, and dense voxel guidance for feature control. However, despite efficiency improvements (Song, Sun, and Yin 2024), video generation models remain computationally intensive (Yang et al. 2024c; Wan et al. 2025) and preclude the scalability necessary for real-time simulation. Furthermore, they struggle with coherent novel-view rendering over long driving trajectories and suffer from a lack of causality. In the top of Tab. 1, we provide a comparison of relevant video synthesis methods for driving scenes.

**3D Generation.** To guarantee consistency and causality, recent approaches separate from 2D image and video generation and focus on the explicit generation of individual, object-centric 3D structures. These methods (Xiang et al. 2024; Li et al. 2025; Gao\* et al. 2024; Sargent et al. 2023; Xie et al. 2024b) generate consistent 3D structure by leveraging well-defined features and latent spaces, and multi-view observations – enabling them to directly perform explicit diffusion of 3D objects. However, these 3D generative models rely on high-quality synthetic 3D data (Deitke et al. 2023; Collins et al. 2022), and most works are thus constrained to the generation of individual objects – generation of large-scale 3D outdoor scenes remains an open challenge.

**Scene Reconstruction and Generation.** For large-scale scene generation (Kim et al. 2023; Yang et al. 2023a; Lee et al. 2024; Wang et al. 2023a; Bahmani et al. 2023), various 3D representations have been proposed: triplanes (Shue et al. 2023; Lee et al. 2024), semantic occupancy grids, bounding boxes, and 3D maps (Yang et al. 2023a; Xu et al. 2023). However, reliance on generation of explicit priors requires expensive annotation data, and these early methods’ generations lack significantly in photorealism and scale. In contrast, satellite imagery has been used for the city-scale 3D generation of urban environments (Lin et al. 2023; Xie et al. 2024a). More recently, hierarchical voxel diffusion methods (Ren et al. 2024a), have used accumulated LiDAR point clouds to supervise the generation of 3D driving scenes or generate realistic LiDAR point clouds (Zyrianov et al. 2024). However, these methods exclusively generate geometry, without the texture or appearance needed for training perception models – see the bottom of Tab. 1.

Method	Unlimited	Compos-	Causal 3D	Real-Time	View	Control		
	Viewpoints	ability	Geometry	Rendering	Extrapolation	Weather	Time	Map
DriveDreamer (2023b)	✗	✗	✗	✗	✗	✓	✓	✗
WonderJourney (2024b)	✗	✗	✗	✗	✗	✓	✓	✗
Streetscapes (2024)	✗	✗	✗	✗	✗	✓	✓	✓
Vista (2024)	✗	✗	✗	✗	✗	✓	✓	✗
MagicDriveDiT (2024b)	✗	✗	✗	✗	✗	✓	✓	✗
WoVoGen (2024a)	✗	✗	✗	✗	✗	✓	✓	✓
WonderWorld (2024a)	✗	✗	✗	✗	✗	✓	✓	✗
NF-LDM (2023)	✗	✗	✗	✓	✗	✓	✓	✓
InfiniCity (2023)	✗	✗	✓	✓	✓	✗	✗	✓
MagicDrive3D (2024)	✗	✓	✓	✓	✓	✓	✓	✓
CityDreamer (2024a)	✗	✗	✓	✓	✓	✗	✗	✓
InfiniCube (2024b)	✗	✓	✓	✓	✓	✓	✓	✓
GEN3C (2025b)	✗	✗	✓	✗	✓	✓	✓	✗
<b>LSD<sup>3D</sup>-3D (Ours)</b>	✓	✓	✓	✓	✓	✓	✓	✓

Table 1: Video and 3D Generation for Driving Scenes. We review the recent body of video (top) and 3D generation work (bottom).

Neural scene reconstruction (Mildenhall et al. 2021) on the other side has shown promising results for the production of photorealistic visual data. 3D Gaussians (Kerbl et al. 2023) have emerged as prime scene representations, able to explicitly model geometry while also allowing for real-time rendering, enabling scalability. Methods such as OmniRe (Chen et al. 2025b), SplatAD (Hess et al. 2025), SCube (Ren et al. 2024b) or STORM (Yang et al. 2025) are capable of reconstructing 3D texture and geometry as 3D Gaussians from real-world driving videos, allowing for the exploration of novel trajectories. Nevertheless, pure reconstruction methods are still fundamentally limited by the availability of real data to reconstruct from. A natural extension of these works has been undertaken by works like WoVoGen (Lu et al. 2024a) and InfiniCube (Lu et al. 2024b), which replace the real data required for scene reconstruction with generated videos conditioned on scene geometry – the latter then fits these videos onto a set of deformable Gaussian Splats, a first approach in 3D grounding. However, this approach is still inherently limited to the generated video trajectory. Furthermore, despite the explicit 3D representation, the lack of causality in a single original video results in visual artifacts and inconsistencies being baked into the 3D representation. In contrast, distillation is a paradigm which has recently emerged in 3D object generation (Poole et al. 2022; Liang et al. 2024) and sparse reconstruction (Wu et al. 2025, 2024b; ?; Fischer et al. 2025) that is centered around the learning of a 3D representation guided by pretrained text-to-image models. This approach makes it possible to synthesize novel 2D views from arbitrary camera positions, while ensuring 3D consistency and visual fidelity throughout the optimization, thereby bypassing the problem of limited 3D data availability. However, generation by distillation works for object-centric views but does not naturally scale to complex textures and large-scale scenes, as is necessary for driving scene generation. Our work investigates how distillation can be extended for large-scale scene generation, allowing for the *distillation* of image priors as part of the 3D

generation – as opposed to explicit *supervision*.

### 3 Geometry-Grounded 3D Generation

Our method generates large-scale driving scenes with 3D-consistent geometry and texture. We provide an overview of the generative process in Fig. 1. In the following, we first describe our large-scale scene representation before introducing the proposed conditional generation process.

#### 3.1 Scene Representation

**Geometric Layout and Background Environment.** The coarse geometric layout of the scene encodes road, rough vegetation, static vehicles, and building facades. The layout is represented by a mesh  $\mathcal{M} = \{\mathbf{F}_1, \dots, \mathbf{F}_N\}$  where each triangular face is defined by three vertices  $[\mathbf{V}_a, \mathbf{V}_b, \mathbf{V}_c]$ .

We model the background texture at infinity with an environment map (Greene 1986). For a given time of day, weather, and seasonal setting, we introduce into our scenes a corresponding background environment in the form of an equirectangular map  $\mathcal{E}$  of size  $H_{env} \times W_{env} \times 3$ , that offers explicit environment lighting control. Queried as a spherical environment map  $f_{env}(\mathbf{d}, \mathcal{E})$ , it returns a color  $\mathbf{c}$  for any given viewing directions  $\mathbf{d} = (\varphi, \eta) \in \mathbb{R}, (0, 2\pi]$ .

**Gaussian Structure and Texture.** On a finer level, we represent detailed geometry and texture as a set of 2D oriented planar Splats introduced by Huang et al.. Each splat  $\theta_k$  is parametrized by its central point  $\mathbf{p}_k \in \mathbb{R}^3$ , two principal tangential vectors  $\mathbf{t} = (\mathbf{t}_u, \mathbf{t}_v)$  that define their orientation, and a variance controlling its scale per axis  $\mathbf{s} = (s_u, s_v)$ . Our complete scene representation  $\Theta$  is defined as the set of all  $K$  individual Gaussian  $\Theta = [\theta_0, \dots, \theta_K]$ . Complex textures are further modeled by the Gaussian appearance  $\mathbf{c}$  (stored as a set of spherical harmonics) and opacity  $\mathbf{o}$ . The rendered foreground is alpha-composited with the environment map rendered at infinity for each pixel’s viewing direction.

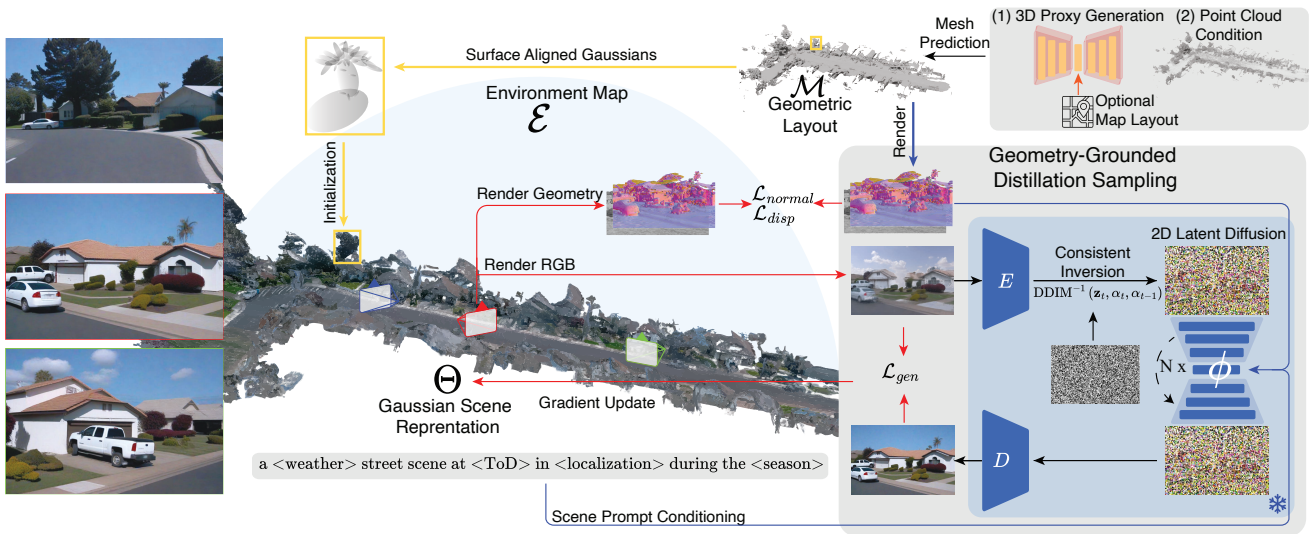


Figure 1: Geometry-Grounded Large-Scale 3D Scene Generation. We generate a large-scale scene as a combination of a coarse geometric layout, an environment map, and a set of Gaussians for texture details, discussed in Sec. 3.1. The geometric layout is either generated, conditioned on a map, or predicted from point-cloud data and guides the overall scene structure. We can further control the setting with a scene prompt, describing time-of-day, season, and weather. Through Geometry-Grounded Distillation Sampling (GGDS), we then further optimize the Gaussian-based scene representation by leveraging 2D priors from the conditional latent diffusion model through consistent diffusion sampling and image-space optimization – together with a set of geometry-grounding regularizers – and generate a causal large-scale scene representation as described in Sec. 3.3.

### 3.2 Geometric Layout Generation

To generate the foreground mesh geometry, we first generate a voxel occupancy  $V$  from a hierarchical latent voxel diffusion model (Rombach et al. 2022), which makes use of a dense, low-resolution and a sparse, high-resolution 3D UNet (Zhou, Du, and Wu 2021; Ren et al. 2024a) as its respective backbones (see the Supplementary Material for details). To enable explicit control of the layout of  $V$ , we condition the diffusion model on a map layout  $M$  to model the conditional probability  $p(V|M)$ . As single-scene generations are limited to  $100m \times 100m$ , we expand the generation size by introducing chunk-wise outpainting. The initial chunk is exclusively conditioned on the locally aligned map  $M_e$ , but each subsequent chunk  $e$  is also conditioned on an overlapping zone with the previous chunk  $e-1$  with  $p(V_e|M_e, V_{e-1})$ . This diffusion model is trained from scratch on aggregated point-clouds and maps of the target street scene dataset. From the generated voxel grid, we then predict the enclosing coarse surface mesh geometry  $\mathcal{M}$  with neural kernel surface reconstruction (Huang et al. 2023).

### 3.3 Geometry-Grounded Scene Generation

Given the generated coarse mesh  $\mathcal{M}$  and a monochromatically initialized environment map  $\mathcal{E}$ , we next generate a textured scene with causal consistency.

**Mesh to Gaussian Representation.** We place Gaussians  $\Theta$  at  $\mathbf{p}_k$ , to represent mesh faces  $\mathbf{F}$  and set orientation, scale, and tangential axes according to the triangle normal  $\mathbf{n}_F$ , and area  $a_F$ , with orientation  $\mathbf{n}_F = \mathbf{t}_u \times \mathbf{t}_v$  and scale  $|\mathbf{s}| = a_F$ .

### Geometry-Grounded Distillation Sampling (GGDS).

We next distill a latent diffusion model (LDM)  $p_{\phi, data}(\mathbf{z})$  on the set of Gaussians  $\Theta$  through a novel iterative optimization method, which we term Geometry-Grounded Distillation Sampling (GGDS). This optimization method is designed to avoid the artifacts that are typically present in existing latent space distillation-based models (Poole et al. 2022; Liang et al. 2024; Hwang et al. 2024) and in ego-centric scenes (see Fig. 3).

In each distillation step, we first obtain an image  $x_i = g(\Theta, \psi_i)$  for viewpoint  $\psi_i$  with the rasterization function  $g$ . We encode the latent  $z_{0,i} = E(x_i)$  from this image and add noise  $\epsilon$  of noise level  $t$  to obtain the noisy latent  $z_t$ . The noise level  $t$  is sampled uniformly between  $t_{max}$  and  $t_{min}$ . The noisy latent  $z_t$  is the denoised for  $N$  steps and decoded, generating a ground-truth image  $\hat{x}_i = D(\hat{z}_{0,i})$  for the respective viewpoint, inducing a loss in image-space. We formulate the objective as the image reconstruction loss between the generated image  $\hat{x}_i$  and the rendered image  $\mathbf{x}_i = g(\Theta, \psi_i)$ :

$$\mathcal{L}_{gen}(\Theta) = \mathbb{E}_{\psi_i, t} [\omega(t) (\|g(\Theta, \psi_i) - \hat{x}_i\| + \mathcal{L}_{LPIPS}(g(\Theta, \psi_i), \hat{x}_i))] \quad (1)$$

where  $\omega(t)$  is the noise-level dependent weight and  $\mathcal{L}_{LPIPS}$  is the perceptual similarity (Zhang et al. 2018). We choose  $N = 5$  independently of  $t$ , which we show allows for higher generation quality for lower noise level  $t$  in the later stages of the scene optimization. To enforce progressive optimization from coarse to fine, the respective generation strength from the image prior is linearly annealed by dropping the lower sampling bound  $t_{min}$ . Directly opti-

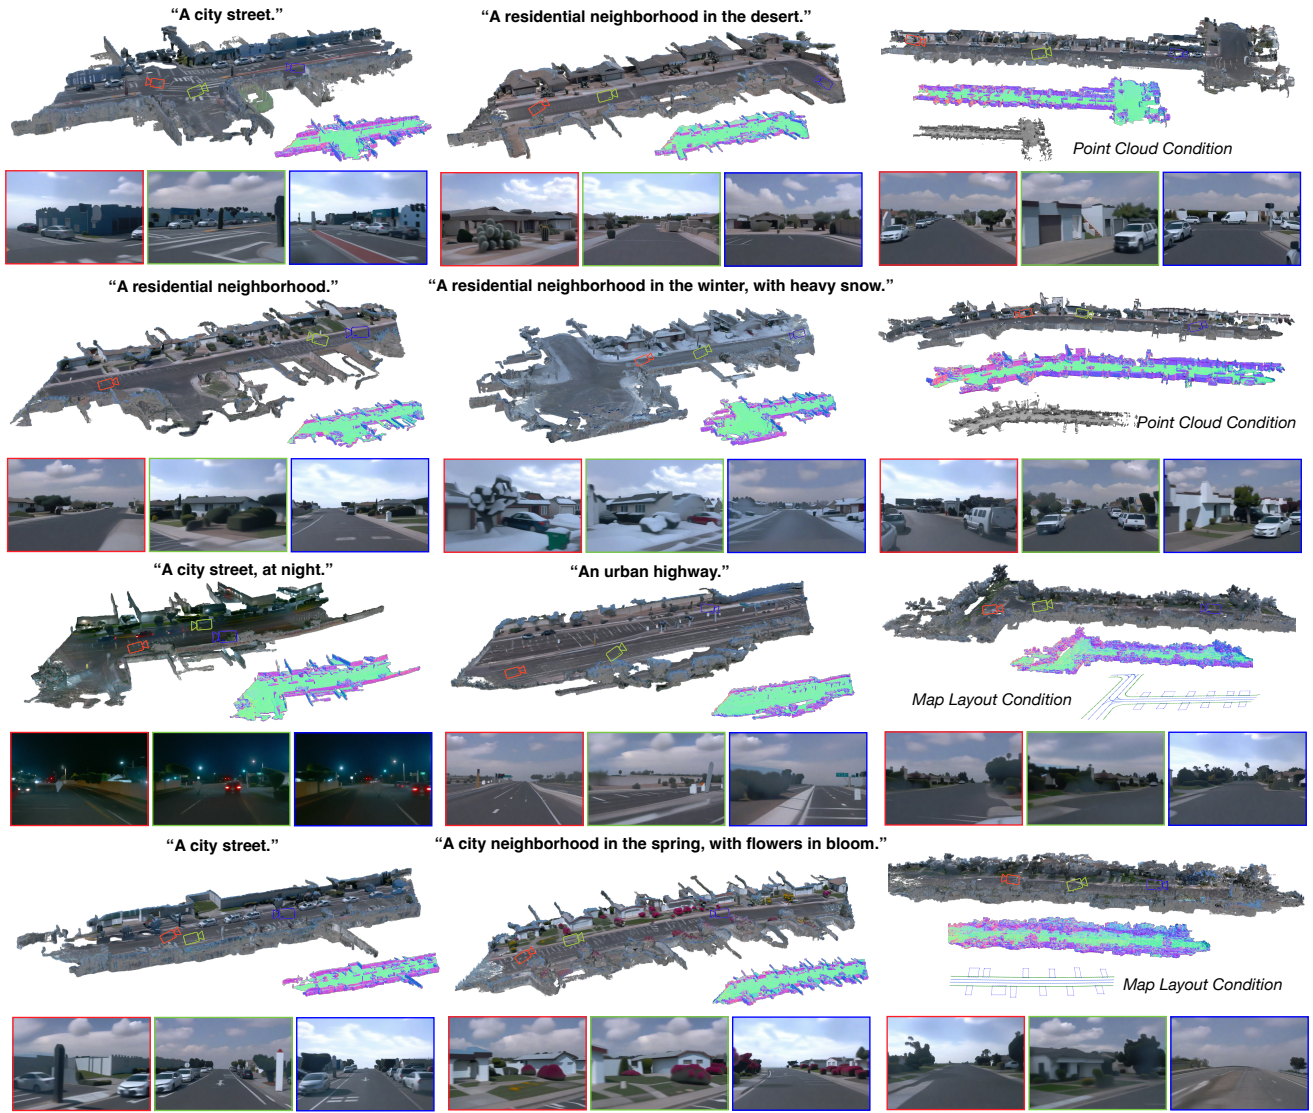


Figure 2: Geometry-Grounded LSD-3D Generations. We visualize 3D scenes generated via our method, alongside the corresponding map of surface normals and selection of novel viewpoints at street level for each of them. In the first two columns, we provide samples of scenes with diversity in time-of-day, season, location, and scene type. In the third column, we provide examples of generated scenes with a point cloud condition and map layout condition..

mizing in image space has significant advantages over score distillation, as our ablation experiments validate, where latent optimization is unable to converge on non-overlapping viewpoints.

To further mitigate randomness that leads to diverging optimization objectives, we enforce consistency between optimization steps at the same viewpoint through DDIM inversion instead of random noise sampling from noise level  $t$ . This ensures a higher level of consistency between the rendered  $\mathbf{x}_i$  and generated images  $\hat{\mathbf{x}}_i$  even in later steps of the optimization, which is in contrast to random sampling, where high noise levels  $t$  can lead to extreme disagreement. We propose a fixed  $N$ -step DDIM inversion (Song, Meng, and Ermon 2020) at any noise-level  $t$  and directly predict

$$\begin{aligned}
 \mathbf{z}_{t,i} &= \text{DDIM}^{-1}(\mathbf{z}_{t-1,i}, \alpha_t, \alpha_{t-1}) \\
 &= \frac{\sqrt{\alpha_t}}{\sqrt{\alpha_t - 1}} \left( \mathbf{z}_{t-1} - \sqrt{1 - \alpha_{t-1}} \epsilon_{\Phi}(\mathbf{z}_{t-1}) \right) \\
 &\quad + \sqrt{1 - \alpha_t} \epsilon_{\Phi}(\mathbf{z}_{t-1}).
 \end{aligned} \tag{2}$$

where  $\mathbf{z}_t$  and  $\mathbf{z}_{t-1}$  represent the noisy latent,  $\epsilon_{\Phi}$  the predicted noise, and  $\{\alpha_t\}_{t=0}^T$  indicate noise level indexing a monotonically increasing time schedule. This allows the model to only introduce changes exactly where needed in each optimization step to satisfy the 2D diffusion prior, also confirmed in Fig. 3. Given this loss objective, we then optimize the Gaussian scene representation through Stochastic Gradient Langevin Dynamics (SGLD) updates (Kherad-



Figure 3: Ablation Experiments. We qualitatively validate the core components of our optimization method. With vanilla SDS, scenes completely fail to converge, necessitating our Gaussian Optimization approach. The proposed texture regularization and initialization approach ensures that scenes converge to a reasonable color distribution, while scenes without them fail. The bottom table reports FID scores without the same components, quantitatively confirming that optimization, geometry-grounding, and texture regularization prove critical for producing high-quality 3D driving scenes.

mand et al. 2024) with

$$\Theta_{k+1} = \Theta_k + \xi (\nabla_{\Theta} \mathcal{L}_{\text{gen}}(\Theta_k)) + \lambda_{\text{noise}} \epsilon, \quad (3)$$

where  $\lambda_{\text{noise}} \epsilon$  is the perturbation of each Gaussian in  $\Theta$ .

As text conditioning  $c$  alone is not a suitable prior to achieve the style of street scenes, we first finetune a LDM  $p_{\phi, \text{fine-tune}}(\mathbf{x})$  on the desired image distribution. Additionally, to avoid scene geometry drifting from the initial mesh, we incorporate disparity conditioning (Zhang, Rao, and Agrawala 2023) on disparity maps computed from the rendered mesh depth  $\mathcal{D} = g_d(\mathcal{M}, \psi_i)$  to the denoising process  $p_{\phi, \text{fine-tune}}(\mathbf{x} | c_{\text{text}}, g_d(\mathcal{M}, \psi_i))$ . This step is crucial for consistency across views and the gradient signal from the 2D diffusion prior on the generated proxy geometry.

**3D Geometry Loss.** Alongside geometry conditioning in the 2D diffusion process, we also regularize all  $\theta_k$  to retain high-quality, smooth 3D geometry. This is achieved by penalizing splat orientation of normal maps  $\mathcal{N}_{\Theta, \psi_i}$  and disparity  $\mathcal{D}_{\Theta, \psi_i}$  from viewpoint  $\psi_i$  with respect to  $\mathcal{N}_{\mathcal{M}, \psi_i}$  and  $\mathcal{D}_{\mathcal{M}, \psi_i}$  rendered from the normals of the proxy geometry as  $\mathcal{L}_{\text{norm}} = \|\mathcal{N}_{\Theta, \psi_i} - \mathcal{N}_{\mathcal{M}, \psi_i}\|$  and  $\mathcal{L}_{\text{disp}} = \|\mathcal{D}_{\Theta, \psi_i} - \mathcal{D}_{\mathcal{M}, \psi_i}\|$ . (4)

Moreover, we apply Gaussian regularization from Huang et al. (2024). Finally, we employ a total-variation (TV) loss on the rendered images to encourage noise reduction in our scenes. We refer to the Supplemental Material for detailed parameter settings and loss composition.

**Deferred Rendering.** We propose a deferred rendering process for novel trajectory videos, inspired by Thies et al. (2019). After scene generation, we use the Gaussian rasterizer to produce an initial frame  $\mathbf{x}_0$  and encode it into a slightly noisy latent  $\mathbf{x}_t$ , which is deferred into final output  $\mathbf{x}_0^{(1)}$  using a fine-tuned T2I model.

This rendering procedure allows us to generate photorealistic images with both low- and high-frequency textures in image space, for instance the road, or a tree, respectively.

## 4 Assessment

In this section, we validate our approach both quantitatively and qualitatively. We conduct an ablation study to validate our design choices and investigate the realism of rendered scenes alongside competing baseline methods.

### 4.1 Ablation Study

We validate the effects of key components of our scene optimization method and visualize the results in Fig 3. Specifically, we analyze the effect of incorporating our improved geometry-grounded distillation sampling (GGDS) approach with SDS, image-space sampling without DDIM inversion, the effectiveness of multi-step denoising and geometry-conditioned diffusion guidance. Without the proposed GGDS, especially in the case of SDS and random noise sampling results, the method experiences catastrophic failure. This highlights the need for two major design choices, which both result in more consistent optimization and distillation of the underlying 3D scene representation. Further multi-step diffusion results in reduced scene generation quality of dark scenes. Without geometry-grounding diffusion guidance, objects in the scene (such as houses or cars) do not follow the conditioning proxy geometry, resulting in inaccurate and flat 3D geometry and poor novel trajectory results. We confirm the visual trend quantitatively in Fig. 3 bottom, which validates that our method (with all proposed components) produces the best visual quality.

### 4.2 Experiments

We validate our method by comparing it against four distinct approaches for 3D scene and video generation.

**Baselines.** We ran two image-to-video generation methods: (a) Vista (Gao et al. 2024), a full driving video model trained on internet-scale driving videos, and (b) WonderWorld (Yu et al. 2024a), in combination with our fine-tuned image model (Podell et al. 2023) on the target dataset (Sun et al. 2020) for fair comparison. For both methods, we fit 2D Gaussian Splats (Huang et al. 2024) to assess geometric causality and novel view synthesis quality.

We also test two conceptually different 3D generation baselines: (c) MagicDrive3D (Gao et al. 2024a), which proposes a driving scene-specific multi-view video diffusion and Gaussian reconstruction pipeline, and (d) GEN3C (Ren et al. 2025b), which fuses geometry prediction from a single image with the Cosmos-Predict (NVIDIA et al. 2025) video diffusion model. Due to the unavailability of public code and models of any candidate (Gao et al. 2024a; Lu et al. 2024b)




Figure 4: Qualitative Comparisons to Video and Scene Generation Methods. Our approach generates an accurate and 3D-consistent scene representation, enabling high-quality novel view synthesis and the generation of unlimited off-trajectory view-points. In contrast, existing baselines WonderWorld (2024a), Vista (2024) combined with Gaussian Splatting (2024; 2023), GEN3C (2025b) and MagicDrive3D (2024a), which generate driving videos and thus lack implicit spatial modeling, fail to generate a consistent and 3D-plausible scene, precluding the production of novel driving trajectories (please zoom for details).

Method	Training / Finetuning	Seen			Novel			CLIP (2021) $\uparrow$
		FID $\downarrow$	FD <sub>DINOv2</sub> $\downarrow$	FVD $\downarrow$	FID $\downarrow$	FD <sub>DINOv2</sub> $\downarrow$	FVD $\downarrow$	
WonderWorld (2024a) + 2DGS	WOD (2020)	130.17	1333.88	1315.53	<u>220.61</u>	<u>1489.51</u>	1424.56	28.88
Vista (2024) + 2DGS	OpenDV (2024a)/WOD	<u>111.93</u>	1510.25	1023.4	242.03	1805.96	<u>1190.76</u>	26.51
MagicDrive3D (2024) <sup>1</sup>	NuScenes (2020)	139.11	1950.67	2285.5	163.73	2004.21	1665.30	21.03
MagicDrive3D (2024) <sup>1</sup> - WOD	WOD <sup>2</sup>	178.71	1965.31	1585.48	186.36	2004.21	1665.30	21.03
GEN3C (2025b) with Cosmos	Proprietary/WOD	<b>97.71</b>	<u>1331.16</u>	<b>962.02</b>	273.27	2237.03	1617.62	26.57
<b>LSD-3D (Ours)</b>	WOD	119.38	<b>1247.19</b>	<u>989.94</u>	<b>119.18</b>	<b>1227.62</b>	<b>974.36</b>	26.03

Table 2: **Quantitative Evaluation of Generation Quality for 3D Scene Generation** with our proposed method and existing approaches. Best results in each visual quality category are in **bold** and the second best underlined. We report CLIP scores on the right, confirming that our generations adhere to prompts on par with existing methods. We apply the same pre-trained T2I model (2023) for all models, and the same 2DGS (2024) pipeline for optimization, to provide fair comparisons. Significantly lower FVD and FID score on novel views confirm temporal and geometric consistency is achieved by our method through geometrically grounded generations.

at the time of submission, we rely on our own implementation (Layout + Geometry Controlled Video Generation) with MagicDrive3D built on top of the latest diffusion transformer version of MagicDriveDiT (Gao et al. 2024b). See the Supplementary Material for details.

**Computational Requirements.** Each scene is generated with 6000 steps, corresponding to an average time of 2 hours on a single NVIDIA H100 GPU. Comparable methods (Liang et al. 2024) that generate single objects, i.e., dramatically smaller scenes, require similar runtime at 5000 steps. Our method outputs frames at 960p resolution at rates higher than 60 fps, providing real-time rendering capabilities. Scenes are initialized with 1.8 to 2.2 million Gaussians, and the maximum is set to 4 million Gaussians.

**Evaluation Metrics.** For all methods, we quantitatively assess the diversity and quality of our results by computing the Fréchet Inception Distance (FID) and the recently established DINOv2 (Oquab et al. 2023) based Fréchet Distance FD<sub>DINOv2</sub> (Stein et al. 2023). For temporal quality, we use the Fréchet Video Distance (FVD), with a subset of the respective training dataset (Sun et al. 2020; Caesar et al. 2020) as reference distribution. Following (Gao et al. 2024a), we evaluate FID and FD<sub>DINOv2</sub> score on generated results from views seen during the Gaussian Splatting optimization (FID *seen*) as well as from novel views sampled at randomly selected distances from the training ones (FID *novel*). Additionally, we evaluate prompt adherence using

the CLIP score (Radford et al. 2021) with the implementation from (Zhengwentai 2023) on 10 common weather, time-of-day, and localization prompts with 3 samples each.

**Quantitative Results.** We evaluate all methods on a set of 40 generated scenes and across ten different scene attributes (time of day, season, weather). Exact prompts are provided in the Supplementary material. As reported in Tab. 2, the quality of rendered images from our model is on par with state-of-the-art 2D generative methods and is capable of generating scenes closely matching in style and content with the source distribution. However, *competing methods cannot generate a 3D-consistent scene*, resulting in novel views with inferior quality – as seen in the resulting high FID and FVD score for novel views all other methods. We also confirm prompt adherence at par with video models, validating that 3D grounding enables fine-grained text control comparable to generalist models (Yu et al. 2024b).

**Qualitative Results.** We also find significant qualitative differences between baselines and our method and baselines in Fig. 4. Baseline image-to-video models produce scene renderings of variable quality. In fact, as views deviate further from the input image, the rendering quality and 3D

<sup>1</sup>No code publicly available (or from authors). Reimplemented with MagicDriveDiT (2024b) and 2DGS (Huang et al. 2024).

<sup>2</sup>We compare scores on the Waymo (Sun et al. 2020) distributions and released video generations. NuScenes (Caesar et al. 2020) results are reported first for completeness of the evaluation.

consistency deteriorate, yielding a Gaussian representation which is inconsistent beyond the original trajectory. This difference is more pronounced when departing from the generated video trajectory - even 2.5D methods fail to produce consistent novel views for unlimited viewpoints. In contrast, our method generates plausible rendering throughout the entire scene, from any realistic viewpoint - without any loss of appearance or geometric quality.

## 5 Conclusion

We introduce, to our knowledge, the first distillation approach to directly generate large-scale explicit 3D driving scenes. To accomplish this, we propose Geometry-Grounded Distillation Sampling (GGDS), which combines controlled proxy mesh generation with a conditional diffusion prior, producing novel and view-consistent Gaussian splatting scenes. Our approach generates completely unseen driving environments controlled by scene descriptions or traffic map layouts. By the design of our method, every scene is generated causally and 3D-consistent, and allows for real-time rendering of physically-grounded videos along novel trajectories. The approach compares favorably against the most successful existing methods – primarily video diffusion approaches – that struggle with view-consistent rendering and causality, and are fundamentally limited to individual trajectories. As a geometry-grounded approach, we hope to integrate the method with driving simulators and extend the domain beyond autonomous driving – ultimately building towards the goal of fully data-driven simulators.

## Acknowledgments

Felix Heide was supported by an NSF CAREER Award (2047359), a Packard Foundation Fellowship, a Sloan Research Fellowship, a Sony Young Faculty Award, a Project X Innovation Award, a Amazon Science Research Award, and a Bosch Research Award.

## References

Anciukevicius, T.; Xu, Z.; Fisher, M.; Henderson, P.; Bilen, H.; Mitra, N. J.; and Guerrero, P. 2022. RenderDiffusion: Image Diffusion for 3D Reconstruction, Inpainting and Generation. *arXiv*.

Bahmani, S.; Skorokhodov, I.; Rong, V.; Wetzstein, G.; Guibas, L.; Wonka, P.; Tulyakov, S.; Park, J. J.; Tagliasacchi, A.; and Lindell, D. B. 2023. 4d-fy: Text-to-4d generation using hybrid score distillation sampling. *arXiv preprint arXiv:2311.17984*.

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*.

Borkman, S.; Crespi, A.; Dhakad, S.; Ganguly, S.; Hogins, J.; Jhang, Y.-C.; Kamalzadeh, M.; Li, B.; Leal, S.; Parisi, P.; et al. 2021. Unity perception: Generate synthetic data for computer vision. *arXiv preprint arXiv:2107.04259*.

Brock, A.; Donahue, J.; and Simonyan, K. 2018. Large scale GAN training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*.

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 *CVPR*.

Chen, A.; Zheng, W.; Wang, Y.; Zhang, X.; Zhan, K.; Jia, P.; Keutzer, K.; and Zhang, S. 2025a. GeoDrive: 3D Geometry-Informed Driving World Model with Precise Action Control. *arXiv:2505.22421*.

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. 2025b. OmniRe: Omni Urban Scene Reconstruction. In *The Thirteenth International Conference on Learning Representations*.

Collins, J.; Goel, S.; Deng, K.; Luthra, A.; Xu, L.; Gundogdu, E.; Zhang, X.; Yago Vicente, T. F.; Dideriksen, T.; Arora, H.; Guillaumin, M.; and Malik, J. 2022. ABO: Dataset and Benchmarks for Real-World 3D Object Understanding. *CVPR*.

Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3213–3223.

Dauner, D.; Hallgarten, M.; Li, T.; Weng, X.; Huang, Z.; Yang, Z.; Li, H.; Gilitschenski, I.; Ivanovic, B.; Pavone, M.; Geiger, A.; and Chitta, K. 2024. NAVSIM: Data-Driven Non-Reactive Autonomous Vehicle Simulation and Benchmarking. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Deitke, M.; Liu, R.; Wallingford, M.; Ngo, H.; Michel, O.; Kusupati, A.; Fan, A.; Laforte, C.; Voleti, V.; Gadre, S. Y.; VanderBilt, E.; Kembhavi, A.; Vondrick, C.; Gkioxari, G.; Ehsani, K.; Schmidt, L.; and Farhadi, A. 2023. Objaverse-XL: A Universe of 10M+ 3D Objects. *arXiv preprint arXiv:2307.05663*.

Deng, B.; Tucker, R.; Li, Z.; Guibas, L.; Snavely, N.; and Wetzstein, G. 2024. Streetscapes: Large-scale Consistent Street View Generation Using Autoregressive Video Diffusion. In *SIGGRAPH 2024 Conference Papers*.

Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.

Dockhorn, T.; Vahdat, A.; and Kreis, K. 2021. Score-based generative modeling with critically-damped langevin diffusion. *arXiv preprint arXiv:2112.07068*.

Dosovitskiy, A.; Ros, G.; Codevilla, F.; Lopez, A.; and Koltun, V. 2017. CARLA: An open urban driving simulator. In *Conference on robot learning*, 1–16. PMLR.

Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12873–12883.

Ettinger, S.; Cheng, S.; Caine, B.; Liu, C.; Zhao, H.; Pradhan, S.; Chai, Y.; Sapp, B.; Qi, C. R.; Zhou, Y.; Yang, Z.; Chouard, A.; Sun, P.; Ngiam, J.; Vasudevan, V.; McCauley, A.; Shlens, J.; and Anguelov, D. 2021. Large Scale Interactive Motion Forecasting for Autonomous Driving: The Waymo Open Motion Dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 9710–9719.

Fischer, T.; Bulò, S. R.; Yang, Y.-H.; Keetha, N. V.; Porzi, L.; Müller, N.; Schwarz, K.; Luiten, J.; Pollefeys, M.; and Kotschieder, P. 2025. FlowR: Flowing from Sparse to Dense 3D Reconstructions. *arXiv preprint arXiv:2504.01647*.

Gao, R.; Chen, K.; Li, Z.; Hong, L.; Li, Z.; and Xu, Q. 2024a. MagicDrive3D: Controllable 3D Generation for Any-View Rendering in Street Scenes. *arXiv preprint arXiv:2405.14475*.

Gao, R.; Chen, K.; Xiao, B.; Hong, L.; Li, Z.; and Xu, Q. 2024b. MagicDriveDiT: High-Resolution Long Video Generation for Autonomous Driving with Adaptive Control. *arXiv:2411.13807*.

- Gao, R.; Chen, K.; Xie, E.; Hong, L.; Li, Z.; Yeung, D.-Y.; and Xu, Q. 2023. Magicdrive: Street view generation with diverse 3d geometry control. *arXiv preprint arXiv:2310.02601*.
- Gao\*, R.; Holynski\*, A.; Henzler, P.; Brussee, A.; Martin-Brualla, R.; Srinivasan, P. P.; Barron, J. T.; and Poole\*, B. 2024. CAT3D: Create Anything in 3D with Multi-View Diffusion Models. *Advances in Neural Information Processing Systems*.
- 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. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, 3354–3361. IEEE.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Greene, N. 1986. Environment mapping and other applications of world projections. *IEEE computer graphics and Applications*, 6(11): 21–29.
- Gulino, C.; Fu, J.; Luo, W.; Tucker, G.; Bronstein, E.; Lu, Y.; Harb, J.; Pan, X.; Wang, Y.; Chen, X.; Co-Reyes, J. D.; Agarwal, R.; Roelofs, R.; Lu, Y.; Montali, N.; Mouglin, P.; Yang, Z.; White, B.; Faust, A.; McAllister, R.; Anguelov, D.; and Sapp, B. 2023. Waymax: An Accelerated, Data-Driven Simulator for Large-Scale Autonomous Driving Research. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*.
- Harvey, W.; Naderiparizi, S.; Masrani, V.; Weilbach, C.; and Wood, F. 2022. Flexible diffusion modeling of long videos. *Advances in Neural Information Processing Systems*, 35: 27953–27965.
- 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 *Proceedings of the Computer Vision and Pattern Recognition Conference*, 11982–11992.
- Ho, J.; Chan, W.; Saharia, C.; Whang, J.; Gao, R.; Gritsenko, A.; Kingma, D. P.; Poole, B.; Norouzi, M.; Fleet, D. J.; et al. 2022a. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; Salimans, T.; Gritsenko, A.; Chan, W.; Norouzi, M.; and Fleet, D. J. 2022b. Video diffusion models. *Advances in Neural Information Processing Systems*, 35: 8633–8646.
- Hong, W.; Ding, M.; Zheng, W.; Liu, X.; and Tang, J. 2022. CogVideo: Large-scale Pretraining for Text-to-Video Generation via Transformers. *arXiv preprint arXiv:2205.15868*.
- Huang, B.; Yu, Z.; Chen, A.; Geiger, A.; and Gao, S. 2024. 2D Gaussian Splatting for Geometrically Accurate Radiance Fields. In *SIGGRAPH 2024 Conference Papers*. Association for Computing Machinery.
- Huang, J.; Gojcic, Z.; Atzmon, M.; Litany, O.; Fidler, S.; and Williams, F. 2023. Neural Kernel Surface Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4369–4379.
- Hwang, S.; Kim, M.-J.; Kang, T.; Kang, J.; and Choo, J. 2024. VEGs: View extrapolation of urban scenes in 3d gaussian splatting using learned priors. In *European Conference on Computer Vision*, 1–18. Springer.
- Jun, H.; and Nichol, A. 2023. Shap-E: Generating Conditional 3D Implicit Functions. *arXiv:2305.02463*.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8110–8119.
- 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).
- Kheradmand, S.; Rebain, D.; Sharma, G.; Sun, W.; Tseng, J.; Isack, H.; Kar, A.; Tagliasacchi, A.; and Yi, K. M. 2024. 3D Gaussian Splatting as Markov Chain Monte Carlo. *arXiv:2404.09591*.
- Kim, S. W.; Brown, B.; Yin, K.; Kreis, K.; Schwarz, K.; Li, D.; Rombach, R.; Torralba, A.; and Fidler, S. 2023. NeuralField-LDM: Scene Generation With Hierarchical Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8496–8506.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Lee, J.; Lee, S.; Jo, C.; Im, W.; Seon, J.; and Yoon, S.-E. 2024. SemCity: Semantic Scene Generation with Triplane Diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Li, H.; Shi, H.; Zhang, W.; Wu, W.; Liao, Y.; Wang, L.; Lee, L.-h.; and Zhou, P. 2024. DreamScene: 3D Gaussian-based Text-to-3D Scene Generation via Formation Pattern Sampling. *arXiv preprint arXiv:2404.03575*.
- Li, Y.; Zou, Z.-X.; Liu, Z.; Wang, D.; Liang, Y.; Yu, Z.; Liu, X.; Guo, Y.-C.; Liang, D.; Ouyang, W.; et al. 2025. TripoSG: High-Fidelity 3D Shape Synthesis using Large-Scale Rectified Flow Models. *arXiv preprint arXiv:2502.06608*.
- Liang, Y.; Yang, X.; Lin, J.; Li, H.; Xu, X.; and Chen, Y. 2024. Lucidreamer: Towards high-fidelity text-to-3d generation via interval score matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6517–6526.
- Liao, Y.; Xie, J.; and Geiger, A. 2022. Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3292–3310.
- Lin, C. H.; Lee, H.-Y.; Menapace, W.; Chai, M.; Siarohin, A.; Yang, M.-H.; and Tulyakov, S. 2023. Infinicity: Infinite-scale city synthesis. In *Proceedings of the IEEE/CVF international conference on computer vision*, 22808–22818.
- Liu, X.; Zhou, C.; and Huang, S. 2024. 3dgs-enhancer: Enhancing unbounded 3d gaussian splatting with view-consistent 2d diffusion priors. *Advances in Neural Information Processing Systems*, 37: 133305–133327.
- Ljungbergh, W.; Tonderski, A.; Johnander, J.; Caesar, H.; Åström, K.; Felsberg, M.; and Petersson, C. 2024. NeuroNCAP: Photorealistic Closed-loop Safety Testing for Autonomous Driving. *arXiv preprint arXiv:2404.07762*.
- Lu, J.; Huang, Z.; Zhang, J.; Yang, Z.; and Zhang, L. 2024a. WoVoGen: World Volume-aware Diffusion for Controllable Multi-camera Driving Scene Generation. In *European Conference on Computer Vision (ECCV)*.
- Lu, Y.; Ren, X.; Yang, J.; Shen, T.; Wu, Z.; Gao, J.; Wang, Y.; Chen, S.; Chen, M.; Fidler, S.; and Huang, J. 2024b. InfiniCube: Unbounded and Controllable Dynamic 3D Driving Scene Generation with World-Guided Video Models. *arXiv:2412.03934*.
- Luo, S.; Tan, Y.; Huang, L.; Li, J.; and Zhao, H. 2023. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*.

- Miao, X.; Duan, H.; Ojha, V.; Song, J.; Shah, T.; Long, Y.; and Ranjan, R. 2024. Dreamer XL: Towards High-Resolution Text-to-3D Generation via Trajectory Score Matching. *arXiv preprint arXiv:2405.11252*.
- 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.
- NVIDIA; Agarwal, N.; Ali, A.; Bala, M.; Balaji, Y.; Barker, E.; Cai, T.; Chattopadhyay, P.; Chen, Y.; Cui, Y.; Ding, Y.; Dworakowski, D.; Fan, J.; Fenzl, M.; Ferroni, F.; Fidler, S.; Fox, D.; Ge, S.; Ge, Y.; Gu, J.; Gururani, S.; He, E.; Huang, J.; Huffman, J.; Jannaty, P.; Jin, J.; Kim, S. W.; Klár, G.; Lam, G.; Lan, S.; Leal-Taixe, L.; Li, A.; Li, Z.; Lin, C.-H.; Lin, T.-Y.; Ling, H.; Liu, M.-Y.; Liu, X.; Luo, A.; Ma, Q.; Mao, H.; Mo, K.; Mousavian, A.; Nah, S.; Niverty, S.; Page, D.; Paschalidou, D.; Patel, Z.; Pavao, L.; Ramezani, M.; Reda, F.; Ren, X.; Sabavat, V. R. N.; Schmerling, E.; Shi, S.; Stefaniak, B.; Tang, S.; Tchapmi, L.; Tredak, P.; Tseng, W.-C.; Varghese, J.; Wang, H.; Wang, H.; Wang, H.; Wang, T.-C.; Wei, F.; Wei, X.; Wu, J. Z.; Xu, J.; Yang, W.; Yen-Chen, L.; Zeng, X.; Zeng, Y.; Zhang, J.; Zhang, Q.; Zhang, Y.; Zhao, Q.; and Zolkowski, A. 2025. Cosmos World Foundation Model Platform for Physical AI.
- Qhaleb, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Ost, J.; Mannan, F.; Thurey, N.; Knodt, J.; and Heide, F. 2021. Neural scene graphs for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2856–2865.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*.
- Poole, B.; Jain, A.; Barron, J. T.; and Mildenhall, B. 2022. DreamFusion: Text-to-3D using 2D Diffusion. *arXiv*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning*.
- Razavi, A.; Van den Oord, A.; and Vinyals, O. 2019. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32.
- Ren, X.; Huang, J.; Zeng, X.; Museth, K.; Fidler, S.; and Williams, F. 2024a. XCube: Large-Scale 3D Generative Modeling using Sparse Voxel Hierarchies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- 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. 2025a. Cosmos-Drive-Dreams: Scalable Synthetic Driving Data Generation with World Foundation Models.
- Ren, X.; Lu, Y.; Liang, H.; Wu, J. Z.; Ling, H.; Chen, M.; Fidler, F.; Sanja and Williams; and Huang, J. 2024b. SCube: Instant Large-Scale Scene Reconstruction using VoxSplats. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Ren, X.; Shen, T.; Huang, J.; Ling, H.; Lu, Y.; Nimier-David, M.; Müller, T.; Keller, A.; Fidler, S.; and Gao, J. 2025b. Gen3c: 3d-informed world-consistent video generation with precise camera control. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 6121–6132.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Rong, G.; Shin, B. H.; Tabatabaee, H.; Lu, Q.; Lemke, S.; Možeiko, M.; Boise, E.; Uhm, G.; Gerow, M.; Mehta, S.; et al. 2020. Lgsvl simulator: A high fidelity simulator for autonomous driving. In *2020 IEEE 23rd International conference on intelligent transportation systems (ITSC)*, 1–6. IEEE.
- Sargent, K.; Li, Z.; Shah, T.; Herrmann, C.; Yu, H.-X.; Zhang, Y.; Chan, E. R.; Lagun, D.; Fei-Fei, L.; Sun, D.; et al. 2023. Zeronvs: Zero-shot 360-degree view synthesis from a single real image. *arXiv preprint arXiv:2310.17994*.
- Sauer, A.; Schwarz, K.; and Geiger, A. 2022. Stylegan-xl: Scaling stylegan to large diverse datasets. In *ACM SIGGRAPH 2022 conference proceedings*, 1–10.
- Shah, S.; Dey, D.; Lovett, C.; and Kapoor, A. 2018. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and Service Robotics: Results of the 11th International Conference*, 621–635. Springer.
- Shriram, J.; Trevithick, A.; Liu, L.; and Ramamoorthi, R. 2024. Realmdreamer: Text-driven 3d scene generation with inpainting and depth diffusion. *arXiv preprint arXiv:2404.07199*.
- Shue, J. R.; Chan, E. R.; Po, R.; Ankner, Z.; Wu, J.; and Wetzstein, G. 2023. 3d neural field generation using triplane diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20875–20886.
- Singer, U.; Polyak, A.; Hayes, T.; Yin, X.; An, J.; Zhang, S.; Hu, Q.; Yang, H.; Ashual, O.; Gafni, O.; et al. 2022. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising Diffusion Implicit Models. *arXiv:2010.02502*.
- Song, Y.; Sun, Z.; and Yin, X. 2024. SDXS: Real-Time One-Step Latent Diffusion Models with Image Conditions. *arXiv preprint arXiv:2403.16627*.
- Stein, G.; Cresswell, J.; Hosseinzadeh, R.; Sui, Y.; Ross, B.; Villecroze, V.; Liu, Z.; Caterini, A. L.; Taylor, E.; and Loaiza-Ganem, G. 2023. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. *Advances in Neural Information Processing Systems*, 36: 3732–3784.
- Sun, P.; Kretschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; Vasudevan, V.; Han, W.; Ngiam, J.; Zhao, H.; Timofeev, A.; Ettinger, S.; Krivokon, M.; Gao, A.; Joshi, A.; Zhang, Y.; Shlens, J.; Chen, Z.; and Anguelov, D. 2020. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Talwar, D.; Guruswamy, S.; Ravipati, N.; and Eirinaki, M. 2020. Evaluating validity of synthetic data in perception tasks for autonomous vehicles. In *2020 IEEE International Conference On Artificial Intelligence Testing (AITest)*, 73–80. IEEE.
- Team, A.; Zhu, H.; Wang, Y.; Zhou, J.; Chang, W.; Zhou, Y.; Li, Z.; Chen, J.; Shen, C.; Pang, J.; et al. 2025. Aether: Geometric-aware unified world modeling. *arXiv preprint arXiv:2503.18945*.
- Team, G. 2024. Mochi 1. <https://github.com/genmoai/models>.
- Team, T. H. 2025. Hunyuan3D 2.0: Scaling Diffusion Models for High Resolution Textured 3D Assets Generation. *arXiv:2501.12202*.
- Thies, J.; Zollhöfer, M.; and Nießner, M. 2019. Deferred neural rendering: Image synthesis using neural textures. *Acm Transactions on Graphics (TOG)*, 38(4): 1–12.

- Tonderski, A.; Lindström, C.; Hess, G.; Ljungbergh, W.; Svensson, L.; and Petersson, C. 2023. NeuRAD: Neural rendering for autonomous driving. *arXiv preprint arXiv:2311.15260*.
- Vahdat, A.; and Kautz, J. 2020. NVAE: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33: 19667–19679.
- Vahdat, A.; Kreis, K.; and Kautz, J. 2021. Score-based generative modeling in latent space. *Advances in neural information processing systems*, 34: 11287–11302.
- Wan, T.; Wang, A.; Ai, B.; Wen, B.; Mao, C.; Xie, C.-W.; Chen, D.; Yu, F.; Zhao, H.; Yang, J.; et al. 2025. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*.
- Wang, P.; Xu, D.; Fan, Z.; Wang, D.; Mohan, S.; Iandola, F.; Ranjan, R.; Li, Y.; Liu, Q.; Wang, Z.; et al. 2023a. Taming Mode Collapse in Score Distillation for Text-to-3D Generation. *arXiv preprint arXiv:2401.00909*.
- Wang, X.; Zhu, Z.; Huang, G.; Chen, X.; and Lu, J. 2023b. Drive-dreamer: Towards real-world-driven world models for autonomous driving. *arXiv preprint arXiv:2309.09777*.
- Williams, F.; Gojcic, Z.; Khamis, S.; Zorin, D.; Bruna, J.; Fidler, S.; and Litany, O. 2021. Neural Fields as Learnable Kernels for 3D Reconstruction. *arXiv:2111.13674*.
- Wu, J. Z.; Zhang, Y.; Turki, H.; Ren, X.; Gao, J.; Shou, M. Z.; Fidler, S.; Gojcic, Z.; and Ling, H. 2025. Difx3d+: Improving 3d reconstructions with single-step diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 26024–26035.
- Wu, R.; Gao, R.; Poole, B.; Trevithick, A.; Zheng, C.; Barron, J. T.; and Holynski, A. 2024a. CAT4D: Create Anything in 4D with Multi-View Video Diffusion Models. *arXiv:2411.18613*.
- Wu, R.; Mildenhall, B.; Henzler, P.; Park, K.; Gao, R.; Watson, D.; Srinivasan, P. P.; Verbin, D.; Barron, J. T.; Poole, B.; et al. 2024b. Reconfusion: 3d reconstruction with diffusion priors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 21551–21561.
- 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.
- Xiang, J.; Lv, Z.; Xu, S.; Deng, Y.; Wang, R.; Zhang, B.; Chen, D.; Tong, X.; and Yang, J. 2024. Structured 3d latents for scalable and versatile 3d generation. *arXiv preprint arXiv:2412.01506*.
- Xie, H.; Chen, Z.; Hong, F.; and Liu, Z. 2024a. Citydreamer: Compositional generative model of unbounded 3d cities. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9666–9675.
- Xie, K.; Lorraine, J.; Cao, T.; Gao, J.; Lucas, J.; Torralba, A.; Fidler, S.; and Zeng, X. 2024b. LATTE3D: Large-scale Amortized Text-To-Enhanced3D Synthesis. *arXiv preprint arXiv:2403.15385*.
- Xu, Y.; Chai, M.; Shi, Z.; Peng, S.; Skorokhodov, I.; Siarohin, A.; Yang, C.; Shen, Y.; Lee, H.-Y.; Zhou, B.; et al. 2023. Discoscene: Spatially disentangled generative radiance fields for controllable 3d-aware scene synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4402–4412.
- Yang, J.; Gao, S.; Qiu, Y.; Chen, L.; Li, T.; Dai, B.; Chitta, K.; Wu, P.; Zeng, J.; Luo, P.; Zhang, J.; Geiger, A.; Qiao, Y.; and Li, H. 2024a. Generalized Predictive Model for Autonomous Driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Yang, J.; Huang, J.; Chen, Y.; Wang, Y.; Li, B.; You, Y.; Igl, M.; Sharma, A.; Karkus, P.; Xu, D.; Ivanovic, B.; Wang, Y.; and Pavone, M. 2025. STORM: Spatio-Temporal Reconstruction Model for Large-scale Outdoor Scenes. *arXiv preprint arXiv:2501.00602*.
- Yang, S.; Hou, L.; Huang, H.; Ma, C.; Wan, P.; Zhang, D.; Chen, X.; and Liao, J. 2024b. Direct-a-Video: Customized Video Generation with User-Directed Camera Movement and Object Motion. *arXiv preprint arXiv:2402.03162*.
- Yang, Y.; Yang, Y.; Guo, H.; Xiong, R.; Wang, Y.; and Liao, Y. 2023a. Urbangiraffe: Representing urban scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9199–9210.
- 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 *CVPR*.
- Yang, Z.; Teng, J.; Zheng, W.; Ding, M.; Huang, S.; Xu, J.; Yang, Y.; Hong, W.; Zhang, X.; Feng, G.; et al. 2024c. CogVideoX: Text-to-Video Diffusion Models with An Expert Transformer. *arXiv preprint arXiv:2408.06072*.
- Yi, T.; Fang, J.; Wang, J.; Wu, G.; Xie, L.; Zhang, X.; Liu, W.; Tian, Q.; and Wang, X. 2024. GaussianDreamer: Fast Generation from Text to 3D Gaussians by Bridging 2D and 3D Diffusion Models. In *CVPR*.
- Yu, H.-X.; Duan, H.; Herrmann, C.; Freeman, W. T.; and Wu, J. 2024a. WonderWorld: Interactive 3D Scene Generation from a Single Image. *arXiv preprint arXiv:2406.09394*.
- Yu, H.-X.; Duan, H.; Hur, J.; Sargent, K.; Rubinstein, M.; Freeman, W. T.; Cole, F.; Sun, D.; Snavely, N.; Wu, J.; et al. 2024b. Wonderjourney: Going from anywhere to everywhere. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6658–6667.
- Zhang, J.; Zhang, Q.; Zhang, L.; Kompella, R. R.; Liu, G.; and Zhou, B. 2024a. Urban Scene Diffusion through Semantic Occupancy Map. *arXiv preprint arXiv:2403.11697*.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3836–3847.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *CVPR*.
- Zhang, S.; Zhang, Y.; Zheng, Q.; Ma, R.; Hua, W.; Bao, H.; Xu, W.; and Zou, C. 2024b. 3D-SceneDreamer: Text-Driven 3D-Consistent Scene Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10170–10180.
- Zhang, Z.; Long, F.; Pan, Y.; Qiu, Z.; Yao, T.; Cao, Y.; and Mei, T. 2024c. TRIP: Temporal Residual Learning with Image Noise Prior for Image-to-Video Diffusion Models. *arXiv preprint arXiv:2403.17005*.
- Zhengwentai, S. 2023. clip-score: CLIP Score for PyTorch. <https://github.com/taited/clip-score>. Version 0.2.1.
- Zhou, L.; Du, Y.; and Wu, J. 2021. 3D Shape Generation and Completion Through Point-Voxel Diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 5826–5835.
- Zyrianov, V.; Che, H.; Liu, Z.; and Wang, S. 2024. LidarDM: Generative LiDAR Simulation in a Generated World. *arXiv preprint arXiv:2404.02903*.